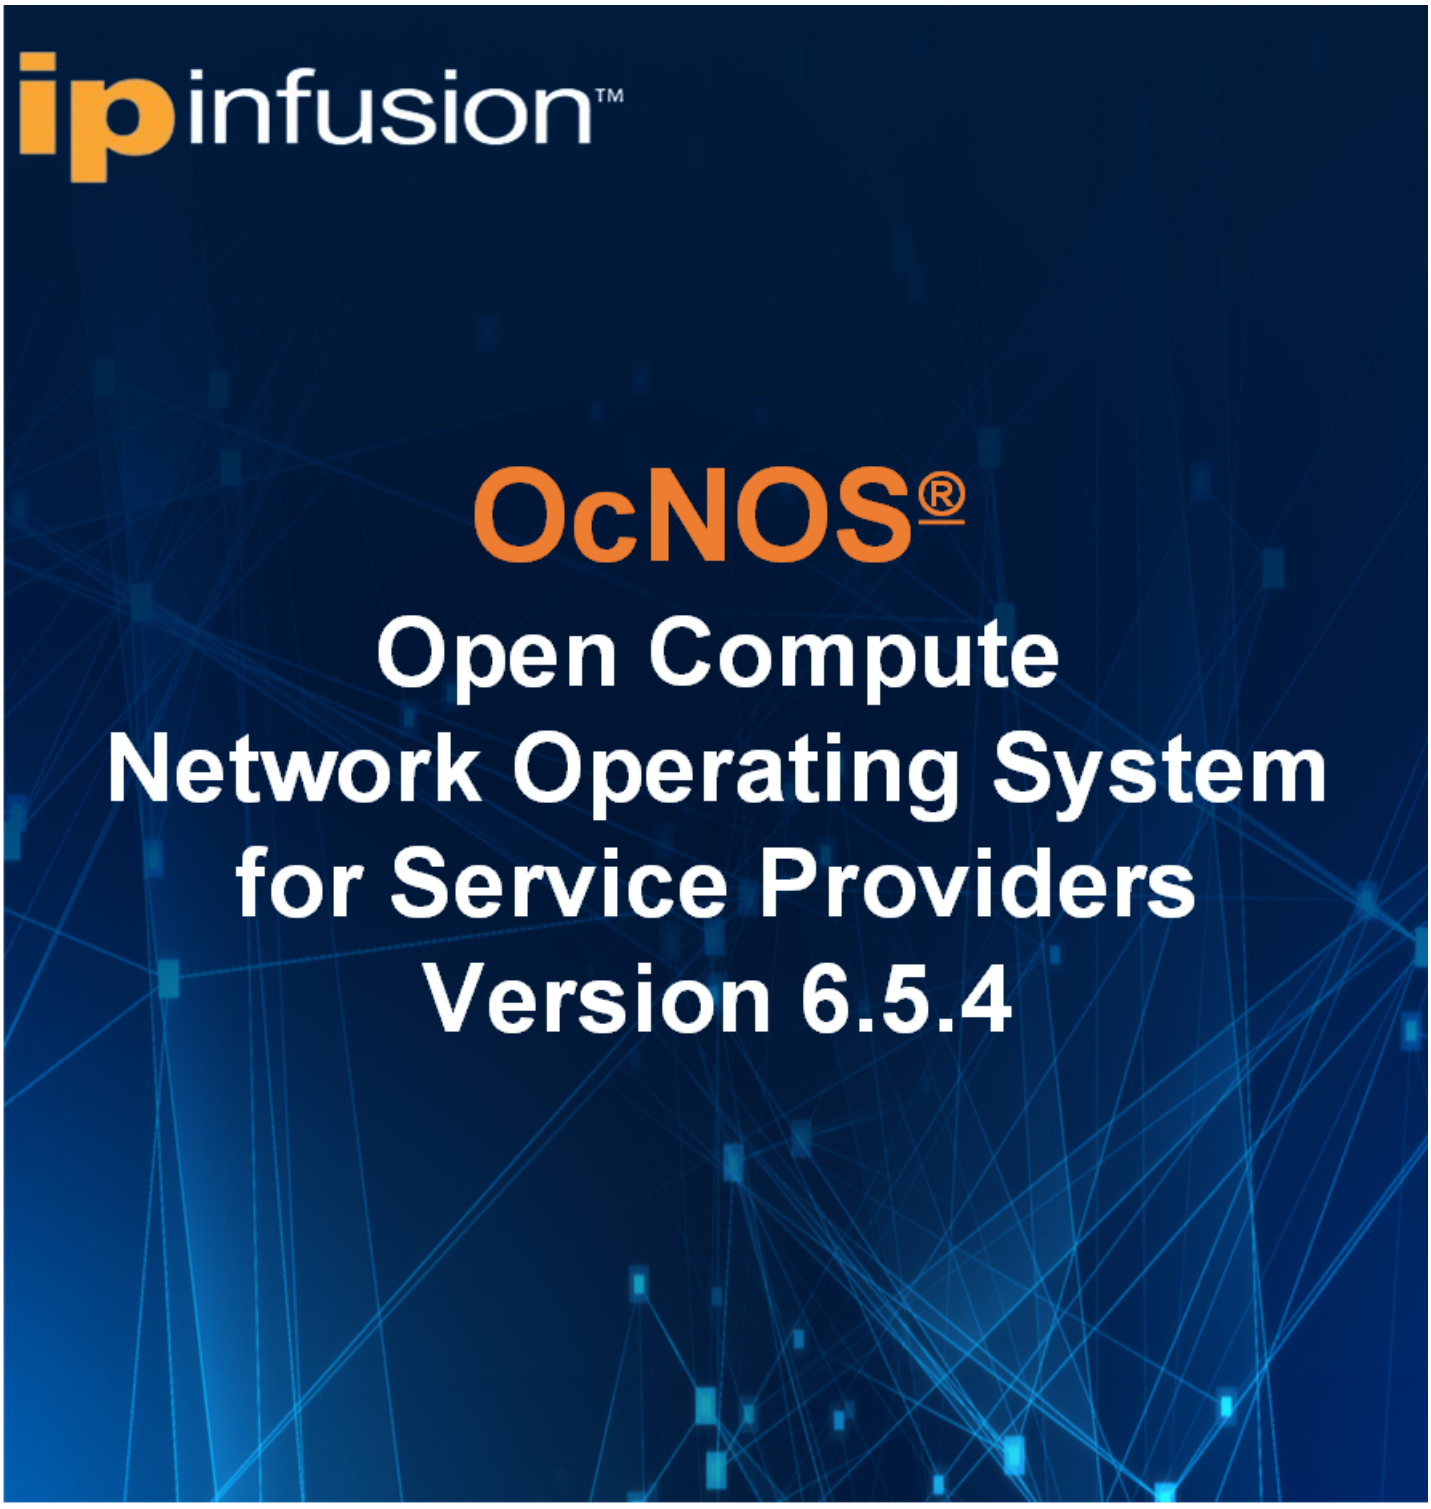

Key Features

April 2025

IP Infusion Inc.

3979 Freedom Circle

Suite 900

Santa Clara, California 95054

http://www.ipinfusion.com/

For support, questions, or comments via E-mail, contact:

support@ipinfusion.com

Contents

Preface

This guide describes how to configure OcNOS.

Audience

This guide is intended for network administrators and other engineering professionals who configure OcNOS.

Conventions

Table 1 on page 10 shows the conventions used in this guide.

Table 1: Conventions

Convention	Description
Italics	Emphasized terms; titles of books
Note:	Special instructions, suggestions, or warnings
`monospaced type`	Code elements such as commands, parameters, files, and directories

IP Infusion Product Release Version

An integer indicates Major, Minor, and Maintenance release versions. Build numbers are for internal tracking and verification of the software build process and are visible to customers as part of the software version number.

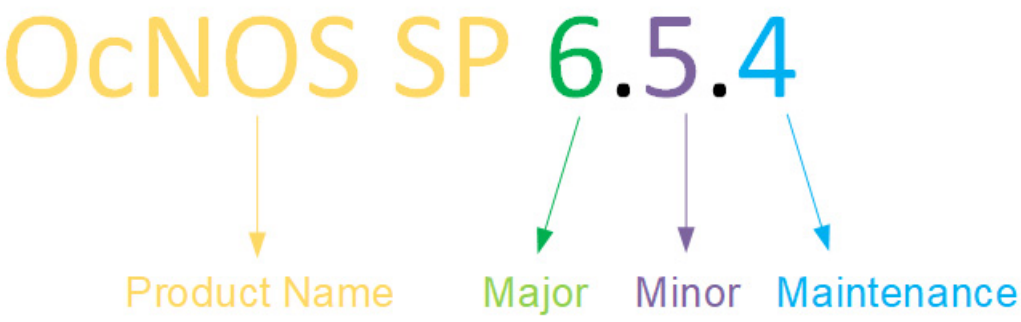

Product Name: IP Infusion Product Family

Major Version: New customer-facing functionality that represents a significant change to the code base; in other words, a significant marketing change or direction in the product.

Minor Version: Enhancements/extensions to existing features, external needs, or internal requirements might be motivated by improvements to satisfy new sales regions or marketing initiatives.

Maintenance Version: It is a collection of product bugs/hotfixes and is usually scheduled every 30 or 60 days, based on the number of hotfixes.

Related Documentation

For information about installing OcNOS, see the *Installation Guide* for your platform.

Feature Availability

The features described in this document that are available depend upon the OcNOS SKU that you purchased. See the *Feature Matrix* for a description of the OcNOS SKUs.

Migration Guide

Check the *Migration Guide* for configuration changes to make when migrating from one version of OcNOS to another.

IP Maestro Support

Monitor devices running OcNOS Release 6.3.4-70 and above using IP Maestro software.

Technical Support

IP Infusion maintains an online technical support site that provides a variety of technical support programs for licensed OcNOS customers at the Technical Assistance Center

Customers and partners enjoy full access to the support website. The site allows customers and partners to open technical support calls, update open calls with new information, and review the status of open or closed calls. The password-protected site includes technical documentation, Release Notes, and descriptions of service offerings.

Technical Documentation

For core commands and configuration procedures, visit: Product Documentation.

For training videos, visit: OcNOS Free Training Videos.

For a list of supported platforms and SKUs of OcNOS features, refer to the OcNOS Feature Matrix.

Technical Sales

Contact the IP Infusion sales representative for more information about the OcNOS solution.

Documentation Disclaimer

The global documentation site is evolving to provide an enhanced website user experience for select topics included in this release. Some guides are now available outside the existing documentation library and can be accessed directly from custom documentation landing pages. These guides offer robust in-built search functionality.

For the latest documentation, visit the product-specific documentation landing page and select the relevant guide.

Comments

If you have comments, or need to report a problem with the content, contact techpubs@ipinfusion.com.

Enhanced Security and Performance

This section describes the network resilience, failover, and error handling enhancements introduced in the Release 6.5.3. No new features are introduced in this section for Release 6.5.3.

Release 6.5.2

- EVPN VXLAN E-Tree
- EVPN MPLS E-Tree
- LDP Tunneling over RSVP-TE
- Hierarchical VPLS
- Auto-Bandwidth with RSVP-TE
- Y.1731 and CFM Over EVPN ELINE Single Home
- Y.1731 and CFM Over EVPN-ELINE Multi-home
- Y.1731 and CFM Over VPLS Sub-Interface
- Y.1731 and CFM Over EVPN ELAN Single Home
- Y.1731 and CFM Over EVPN-ELAN Multi-home
- Y.1731 and CFM Over VPWS Sub-interface

CHAPTER 1 EVPN VXLAN E-Tree

Overview

Ethernet VPN Ethernet-Tree (EVPN E-Tree), is a networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, especially in scenarios where specific services or devices need controlled communication. It categorizes network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Feature Characteristics

Implemented Scenario 1 of the EVPN E-Tree solution, as defined by RFC-8317, designates each Provider Edge (PE) node as either a Leaf or a Root site per Virtual Private Network (VPN) for VXLAN and MPLS EVPN in OcNOS.

Scenario 1: Leaf or Root Site(s) per PE

Scenario 1 involves a topology with three PE nodes: PE-1, PE-2, and PE-3. PE-1 and PE-2 are Multi-Homed nodes (MH-1 and MH-2), with PE-3 acting as the Root node. PE-1 and PE-2 function as Leaf nodes and are part of a single home access interface (SH1 and SH2).

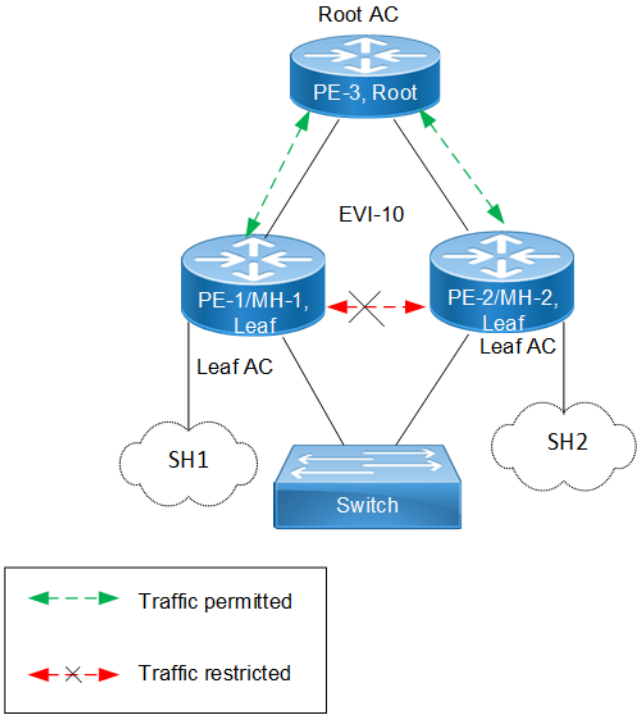

EVPN E-Tree

The classification ensures that communication follows specific rules:

- Communication between Leaf hosts is restricted, as indicated by red dotted lines with a cross mark (X) in the topology diagram. However, communication between Leaf and Root nodes, as well as between Root nodes, is permitted, marked by green dotted lines.

- Leaf nodes within PE-1 and PE-2 are isolated from each other, preventing intra-PE communication.

The scenario 1 is achieved through two main concepts:

1. **Inter-PE Communication**
   - The inter-PE Route Target (RT) Constraint Method is applicable only to Single-Homing (SH) devices. Two RTs per broadcast domain are utilized, with Leaf PEs exporting Leaf RTs and Root nodes exporting Root RTs. Leaf nodes import only Root RTs, allowing communication with Root PEs while preventing communication with other Leaf nodes. RT constraints limit the import of specific EVPN routes (MAC-IP and IMET routes) to designated paths for inter-PE communication.
   - IPI employs a proprietary method to support inter-PE connectivity for both SH and MH devices, using BGP extended community to advertise Leaf Indication in BGP routes and influence traffic flow for both Unicast and BUM traffic. This method enables implementation of ARP or ND cache suppression and MAC mobility sub-features specified in RFC-7432.
2. **Intra-PE communication:** Local Split Horizon controls intra-PE communication between Attachment Circuits (ACs) within Leaf PE nodes, ensuring that traffic between ACs does not egress to other Leaf ACs.

   Note: This functionality depends on hardware capabilities.

Benefits

EVPN E-Tree offers benefits in networking environments by providing efficient traffic control, enhanced security, scalability, and improved performance.

Efficient Traffic Control: EVPN E-Tree allows for efficient control over traffic within network broadcast domains. By segregating nodes into Leaf and Root categories, it enables precise management of communication flows, ensuring the traffic is directed only where needed.

Enhanced Security: The isolation of Leaf hosts from each other adds a layer of security to the network. This prevents unauthorized communication between devices within the same broadcast domain, reducing the risk of data breaches and unauthorized access.

Scalability: EVPN E-Tree is scalable, making it suitable for networks of various sizes and complexities. Whether deploying in small-scale environments or large enterprise networks, EVPN E-Tree offers flexibility and scalability to meet evolving business needs.

Improved Performance: By controlling communication paths and optimizing traffic flows, EVPN E-Tree can improve network performance. This ensures that critical data packets are delivered efficiently, reducing latency and enhancing overall network performance.

Prerequisites

In setting up a VXLAN EVPN network, certain prerequisites are essential to ensure proper functionality and connectivity.

Ensure VXLAN EVPN Configuration: Confirm that VXLAN, EVPN VXLAN, and VXLAN filtering are already enabled in the network as they are required for VXLAN EVPN Multihoming.

Define Interfaces and Loopback Addresses: Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific system MAC addresses (Ethernet Segment Identifier (ESI) values) for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

Configure OSPF and BGP for Dynamic Routing: Enable OSPF to facilitate dynamic routing within the network. Define OSPF router IDs to match loopback IP addresses and add network segments to OSPF areas for proper route

distribution. Additionally, establish BGP sessions to advertise routes between different nodes. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

Leaf Node

1. Enable VXLAN and EVPN MH

   Enable features like VXLAN and EVPN Multihoming, VXLAN filtering, and quality of service (QoS) capabilities on all Leaf nodes.

   ```
   !
   nvo vxlan enable
   !
   evpn vxlan multihoming enable
   !
   qos enable
   !
   ```

2. Configure Interfaces and Loopback

   Define a port channel interface (`po1`) as an L2 interface and assign the system MAC (`0000.0000.1111`) as the ESI value. Designate an interface (`xe7`) as a member port of `po1`. Assign the loopback IP address (`1.1.1.1`) to Leaf node, and set IP addresses (`10.10.10.1` and `10.10.11.1`) to interfaces (`xe45` and `xe49/2`), respectively, for connectivity with Spine nodes.

   ```
   !
   interface po1
    switchport
    evpn multi-homed system-mac 0000.0000.1111
   !
   interface lo
    ip address 1.1.1.1/32 secondary
   !
   interface xe7
    channel-group 1 mode active
   !
   interface xe45
    ip address 10.10.10.1/24
   !
   interface xe49/2
    ip address 10.10.11.1/24
    exit
   !
   ```

3. Configure OSPF

   In OSPF router mode, set the router ID (`1.1.1.1`), to match the loopback IP address. Add the loopback network (`1.1.1.1/32`) and networks (`10.10.10.0/24 and 10.10.11.0/24`) connected to Spine nodes in OSPF area `0`. Enable Bidirectional Forwarding Detection (BFD) on all OSPF interfaces for faster convergence.

   ```
   !
   router ospf 100
    ospf router-id 1.1.1.1
    bfd all-interfaces
    network 1.1.1.1/32 area 0.0.0.0
    network 10.10.10.0/24 area 0.0.0.0
    network 10.10.11.0/24 area 0.0.0.0
   !
   ```

4. Configure BGP

   In BGP router mode, set the router ID (`1.1.1.1`) to match the loopback IP address. Specify the loopback IP address of each Leaf node as neighbors with their respective remote AS numbers. Configure the loopback as the update source for each neighbor and set the advertisement interval (`0`) for rapid convergence. In L2VPN EVPN address family mode, activate each Leaf node (`2.2.2.2, 3.3.3.3, 4.4.4.4`) to establish connections within the EVPN address family.

```
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 2.2.2.2 activate
 neighbor 3.3.3.3 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
```

5. Configure VRF

   In VRF mode, create a MAC routing or forwarding instance (`VRF1`). Assign the Route Distinguisher (RD) value (`1.1.1.1:100`) and set both import and export route-target value (`100:100`) . Ensure that the same route-target value is configured on all Leaf nodes for MAC VRF to maintain consistency.

```
!
mac vrf VRF1
 rd 1.1.1.1:100
 route-target both 100:100
!
```

Spine Node

1. Configure Interfaces and Loopback

   Enable QoS and assign specific IP addresses to loopback interfaces. Configure IP addresses for interfaces connected to each Leaf node.

```
!
qos enable
!
interface ce1/2
 ip address 40.40.40.2/24
!
interface ce1/4
 ip address 10.10.10.2/24
!
interface ce24/1
 ip address 30.30.30.2/24
```

```
!
interface ce27/1
 ip address 20.20.20.2/24
!
interface lo
 ip address 5.5.5.5/32 secondary
!
```

2. Configure OSPF

   In OSPF router mode, set the router ID (`5.5.5.5`), to match the loopback IP address. Add the loopback network (`5.5.5.5/32`) and networks (`10.10.10.0/24, 20.20.20.0/24, 30.30.30.0/24, and 40.40.40.0/24`) connected to Leaf nodes in OSPF area `0`. Enable BFD on all OSPF interfaces for faster convergence.

```
!
router ospf 100
 ospf router-id 5.5.5.5
 bfd all-interfaces
 network 5.5.5.5/32 area 0.0.0.0
 network 10.10.10.0/24 area 0.0.0.0
 network 20.20.20.0/24 area 0.0.0.0
 network 30.30.30.0/24 area 0.0.0.0
 network 40.40.40.0/24 area 0.0.0.0
!
```

Configure Switch

Set up an IEEE VLAN bridge, enabling VLANs and associating them with bridge `1`. Configure interfaces (`xe57, po1, xe46, xe47`) to be part of bridge `1`, setting them as hybrid ports with VLAN (`1000`) allowed and egress-tagged enabled. Designate interfaces connected to Leaf nodes (`xe46 and xe47`) as member ports of `po1`.

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
 vlan-reservation 4000-4094
 vlan 1000 bridge 1 state enable
!
interface po1
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe46
 channel-group 1 mode active
!
interface xe47
 channel-group 1 mode active
!
interface xe57
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
```

Configuration

Configure various nodes within the topology to set up a VXLAN EVPN E-Tree network.

Topology

The sample topology includes Leaf Nodes (VTEP1, VTEP2, VTEP3, and VTEP4), Spine Nodes (SPINE1 and SPINE2), and Switches (SWITCH1 and SWITCH2).

VTEP1 and VTEP2 belong to Multi-homed group 1 (MH1) with po1, while VTEP3 and VTEP4 are in Multi-homed group 2 (MH2) with po2. VTEP2 and VTEP4 connect to single home access ports SH1 and SH2, respectively. All VTEPs link to Spine nodes SPINE1 and SPINE2. SWITCH1 is multi-homed to VTEP1 and VTEP2, and SWITCH2 connects to VTEP3 and VTEP4.

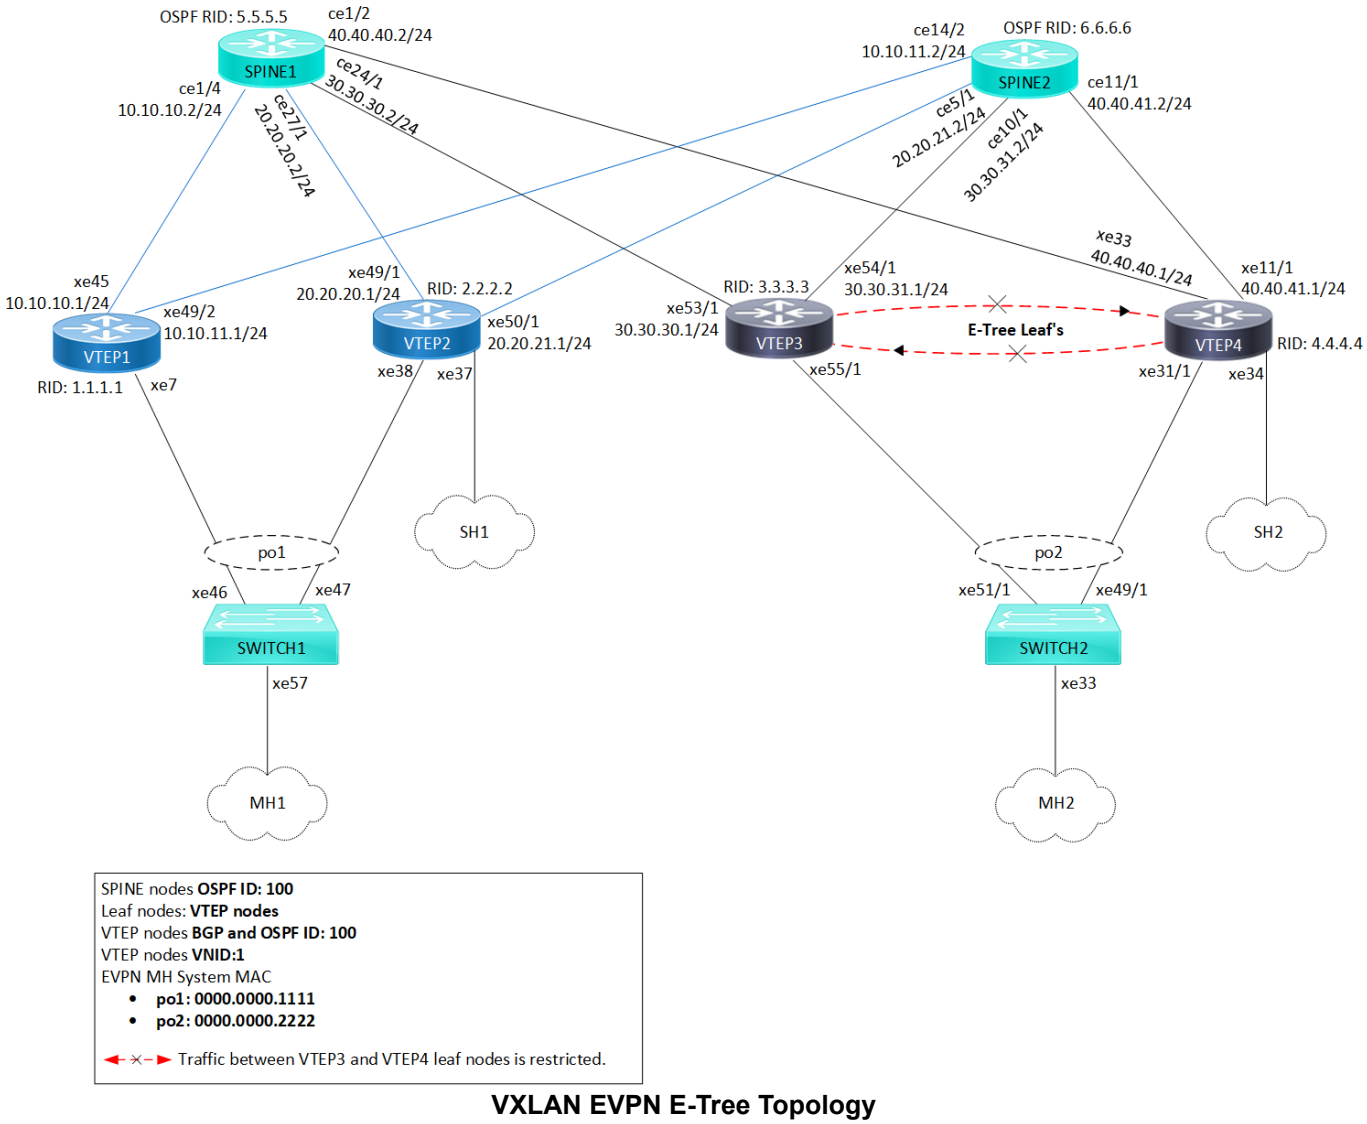

VXLAN EVPN E-Tree Topology

Note: Before configuring E-Tree, meet all Prerequisites for the following nodes:

- Leaf nodes: VTEP1, VTEP2, VTEP3, and VTEP4

- Spine nodes: SPINE1 and SPINE2
- Switches: SWITCH1 and SWITCH2

Enable EVPN E-Tree

The following E-Tree configurations applies to the VTEP nodes within the VXLAN network.

1. Enable EVPN E-Tree on VTEP3 and VTEP4 nodes, allowing them to participate in E-Tree functionality within the VXLAN network, controlling traffic and establishing hierarchical connections between Leaf nodes in the network architecture.

   ```
   (config)#evpn etree enable
   ```

2. Set the ESI hold time (`90 seconds`) on all VTEP nodes to allow the tunnel to establish during VXLAN initialization before bringing up the ESI. Configure the source VTEP IP address (`3.3.3.3`) which serves as the global identifier for VXLAN encapsulation and decapsulation within the network, facilitating proper communication and tunnel establishment.

   ```
   (config)#evpn esi hold-time 90
   (config)#nvo vxlan vtep-ip-global 3.3.3.3
   ```

3. Define VXLAN identifier (`10`) with ingress replication and disabled inner VLAN ID (VID) for **E-Tree leaf nodes** (VTEP3 and VTEP4) to support hierarchical connectivity and traffic control within the VXLAN network. This configuration allows for efficient replication of traffic at the ingress point and ensures that inner VLAN IDs are disabled, optimizing the functionality of E-Tree leaf nodes within the network architecture. On the VXLAN tenant node, assign VRF (`VRF1`) to EVPN-BGP for carrying EVPN routes within the VXLAN network.

   ```
   (config)#nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
   (config-nvo)#vxlan host-reachability- protocol evpn-bgp VRF1
   (config-nvo)#exit
   ```

4. Enable port-VLAN mapping (`po2`) with VLAN ID (`1000`) to facilitate multi-homed access on all VTEP nodes. Map VXLAN identifier (`10`) to the access port for VXLAN connectivity.

   ```
   (config)#nvo vxlan access-if port-vlan po2 1000
   (config-nvo-acc-if)#map vnid 10
   (config-nvo-acc-if)#exit
   (config)#commit
   ```

Validation

Use the show commands described in this section to verify the network for proper VXLAN EVPN E-Tree configuration.

Verify OSPF sessions between the VTEP nodes and the SPINEs within the VXLAN network using the `show ip ospf neighbor` command. This command displays OSPF neighbor details, including the state of the OSPF neighbor relationship. A `State` of `Full/DR` indicates a fully adjacent and operational state between the routers, confirming proper OSPF connectivity within the network.

```
VTEP1#show ip ospf  neighbor

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID     Pri   State    Dead Time   Address      Interface           Instance ID
5.5.5.5           1   Full/DR  00:00:32    10.10.10.2      xe45                   0
6.6.6.6           1    Full/DR  00:00:30   10.10.11.2      xe49/2                 0
```

Verify the BGP session status between VTEPs, using the `show bgp l2vpn evpn summary` command output. The `Up/Down` field indicates the duration for which the BGP session has been up or down.

```
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 9
```

```
1 BGP AS-PATH entries
0 BGP community entries

Neighbor    V    AS    MsgRcv     MsgSen  TblVer  InQ   OutQ Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2.2.2.2     4   100    34         28       7      0      0   00:07:37  9                3      4      1      1         0
3.3.3.3     4   100    30         33       8      0      0   00:07:34  6                3      2      1      0         0
4.4.4.4     4   100    31         28       7      0      0   00:07:37  8                3      4      1      0         0

Total number of neighbors 3

Total number of Established sessions 3
```

To validate the BGP L2VPN output on VTEPs and check MAC-IP routes and ESI information, use the `show bgp l2vpn evpn` command output. This command verifies routes with status code `i (internal) and EVPN route types 2 and 4`, displaying detailed information for each VTEP nodes.

```
VTEP1#show bgp l2vpn evpn
BGP table version is 9, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer          Encap

RD[1.1.1.1:100] VRF[VRF1]:
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*> [2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1000]:[32,100.100.100.1]:[10]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*> [2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1001]:[128,1000::1][10]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i                   2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[2]:[0]:[10]:[48,0000:2000:2000]:[32,200.200.200.1]:[10]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[2]:[0]:[10]:[48,0000:2000:2001]:[128,2000::1][10]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
* i                   4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[2]:[0]:[10]:[48,0000:4000:4000]:[32,104.104.104.1]:[10]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
* i[2]:[0]:[10]:[48,0000:4000:4001]:[128,1004::1][10]
                       4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*> [3]:[10]:[32,1.1.1.1]
                       1.1.1.1              0        100       32768  i  ----------      VXLAN
* i[3]:[10]:[32,2.2.2.2]
                       2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
* i[3]:[10]:[32,3.3.3.3]
                       3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
```

```
* i[3]:[10]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[1.1.1.1:64512] VRF[evpn-gvrf-1]:
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
*> [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]
                      1.1.1.1              0        100       32768  i  ----------      VXLAN
* i[4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:100]
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1000]:[32,100.100.100.1]:[10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1001]:[128,1000::1][10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[0]:[10]:[48,0000:2000:2000]:[32,200.200.200.1]:[10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[2]:[0]:[10]:[48,0000:2000:2001]:[128,2000::1][10]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[3]:[10]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[2.2.2.2:64512]
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN
*>i[4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2         VXLAN

RD[3.3.3.3:100]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN
*>i[3]:[10]:[32,3.3.3.3]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN

RD[3.3.3.3:64512]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      3.3.3.3              0        100       0    i  3.3.3.3         VXLAN

RD[4.4.4.4:100]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[10]:[48,0000:4000:4000]:[32,104.104.104.1]:[10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[2]:[0]:[10]:[48,0000:4000:4001]:[128,1004::1][10]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN
*>i[3]:[10]:[32,4.4.4.4]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

RD[4.4.4.4:64512]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
                      4.4.4.4              0        100       0    i  4.4.4.4         VXLAN

Total number of prefixes 42
```

Validate the LAG interfaces (`po1` and `po2`) are up for MH1 and MH2 by reviewing the `show etherchannel summary` output. Check the `Link` and `sync` fields, where link displays the port channel interface and ID number, and sync indicates whether MAC address synchronization is enabled to forward Layer 3 packets arriving on these interfaces.

```
VTEP1#show etherchannel summary
  Aggregator po1 100001
  Aggregator Type: Layer2
  Admin Key: 0001 - Oper Key 0001
     Link: xe7 (5005) sync: 1
```

Validate the status of NVO VXLAN on VTEPs by examining the output of the `show nvo vxlan` command. The `DF-Status` field displays the forwarding status of VXLAN tunnels as a Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF).

```
VTEP1#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      1.1.1.1          4.4.4.4
10       ----         L2       NW   ----      ----                            ----      ----      1.1.1.1          3.3.3.3
10       ----         L2       NW   ----      ----                            ----      ----      1.1.1.1          2.2.2.2
10       ----         --       AC   po1       00:00:00:00:00:11:11:00:00:00 1000      DF        ----             ----

Total number of entries are 4

VTEP2#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      2.2.2.2          4.4.4.4
10       ----         L2       NW   ----      ----                            ----      ----      2.2.2.2          1.1.1.1
10       ----         L2       NW   ----      ----                            ----      ----      2.2.2.2          3.3.3.3
10       ----         --       AC   xe37      --- Single Homed Port ---     1000      ----      ----             ----
10       ----         --       AC   po1       00:00:00:00:00:11:11:00:00:00 1000      NON-DF    ----             ----

Total number of entries are 5

VTEP3#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          2.2.2.2
10       ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          1.1.1.1
10       ----         L2       NW   ----      ----                            ----      ----      3.3.3.3          4.4.4.4
10       ----         --       AC   po2       00:00:00:00:00:22:22:00:00:00 1000      DF        ----             ----

Total number of entries are 4

VTEP4#show nvo vxlan
VXLAN Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged
```

```
VNID     VNI-Name     VNI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
10       ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          2.2.2.2
10       ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          3.3.3.3
10       ----         L2       NW   ----      ----                            ----      ----      4.4.4.4          1.1.1.1
10       ----         --       AC   xe34      --- Single Homed Port ---       1000      ----      ----             ----
10       ----         --       AC   po2       00:00:00:00:00:22:22:00:00:00   1000      NON-DF    ----             ----

Total number of entries are 5
```

Validate the NVO VXLAN tunnel status on VTEPs by reviewing the output of the `show nvo vxlan tunnel` command. The `Status` field indicates the current status of each tunnel. In this case, all three tunnels between VTEPs and their respective destinations are marked as `Installed`, confirming that these tunnels are successfully established and operating.

```
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source            Destination      Status        Up/Down       Update
================================================================================
1.1.1.1           4.4.4.4          Installed     00:02:26      00:01:58
1.1.1.1           3.3.3.3          Installed     00:02:26      00:01:55
1.1.1.1           2.2.2.2          Installed     00:02:25      00:01:55

Total number of entries are 3
```

Validate the VXLAN access interface status on VTEPs by examining the output of the `show nvo vxlan access-if brief` command. The `up admin and link status` confirms that the access port associated with VXLAN is active and functioning properly on the VTEP nodes.

```
VTEP1#show nvo vxlan access-if brief

                      Inner                          Admin    Link
Interface  Vlan       vlan   Ifindex    Vnid        status   status
-------------------------------------------------------------------
po1        1000       ---    0x7a120    10          up       up

Total number of entries are 1
```

Static MAC-IP Advertisement

Configure static MAC-IP advertisement through SH and MH VTEPs from Root and Leaf nodes. Advertise static MAC addresses for IPv4 and IPv6 from MH1, MH2, SH1, and SH2 VTEPs. Ensure that VTEP1 and VTEP2 in MH1 have the same MAC addresses configured under the port-channel access port. Symmetrical configurations between MH VTEPs should be maintained.

Configure MH1 and MH2 VTEPs

Configure static MAC addresses for IPv4 (`100.100.100.1`) and IPv6 (`1000::1`) under the VXLAN MH access-port (`po1`) with VLAN ID (`1000`). Ensure that identical MAC addresses are set up within the MH1-VTEPs for advertisement. Apply similar configurations to MH2-VTEPs for static MAC-IP advertisement.

```
!
nvo vxlan access-if port-vlan po1 1000
 map vnid 10
 mac 0000.1000.1000 ip 100.100.100.1
 mac 0000.1000.1001 ipv6 1000::1
!
```

Configure SH1 and SH2 VTEPs

Configure static MAC addresses for IPv4 (`200.200.200.1`) and IPv6 (`2000::1`) under the VXLAN SH access-port (`xe37`) with VLAN ID (`1000`) on SH1 (VTEP2). This setup ensures that SH1 advertises these static MAC addresses over the specified VXLAN access-port. Repeat similar configurations for SH2 (VTEP4) using different static MAC addresses for both IPv4 and IPv6.

```
!
nvo vxlan access-if port-vlan xe37 1000
 map vnid 10
 mac 0000.2000.2000 ip 200.200.200.1
 mac 0000.2000.2001 ipv6 2000::1
!
```

Validation

Verify the MAC table entries on MH VTEPs (MH1 and MH2) and the SH VTEPs (VTEP2 and VTEP4). The MAC addresses are advertised using the ESI values from VTEP1 and VTEP2 for MH1, and from VTEP3 and VTEP4 for MH2. Additionally, verify the VTEP IP addresses associated with SH VTEP2 and VTEP4 for MAC advertisement.

In the output of the `show nvo vxlan mac-table` command on all VTEP nodes, the MAC entries advertised from Leaf VTEPs will have the `LeafFlag` field status `set`.

Note:

- MAC IPv4 or IPv6 configured under SH Leaf VTEP access port will be advertised to the Root VTEP and other Leaf VTEPs.
- MAC IPv4 or IPv6 configured under an MH Leaf VTEP access port must be symmetric and will be advertised to both the Root VTEP and other leaf VTEPs.
- MAC IPv4 or IPv6 configured under either SH or MH Root VTEP will be advertised to both the Root VTEP and the Leaf VTEPs.
- The Leaf-to-Leaf communication will display MAC status and tunnel status per VNI as Leaf type. The MAC will be in the discard state in the BCM shell.

```
VTEP1#show nvo vxlan mac-table
=============================================================================================================================================
                                                VXLAN MAC Entries
=============================================================================================================================================
VNID Interface VlanId  In-VlanId Mac-Addr       VTEP-Ip/ESI                    Type   Status  MAC move AccessPortDesc LeafFlag
---------------------------------------------------------------------------------------------------------------------------------------------
10   po1       1000    ----     0000.1000.1000  00:00:00:00:00:11:11:00:00:00 Static Local  ------- 0 -------         ----
10   po1       1000    ----     0000.1000.1001 00:00:00:00:00:11:11:00:00:00  Static Local ------- 0 -------          ----
10   ----      ----    ----     0000.2000.2000 2.2.2.2                        Static Remote   -------  0  -------     ----
10   ----      ----    ----     0000.2000.2001 2.2.2.2                        Static Remote   -------    0  -------   ----
10   ----      ----    ----     0000.3000.3000 00:00:00:00:00:22:22:00:00:00  Static Remote   -------    0  -------   set
10   ----      ----    ----     0000.3000.3001 00:00:00:00:00:22:22:00:00:00  Static Remote   -------    0   -------  set
10   ----      ----    ----     0000.4000.4000 4.4.4.4                        Static Remote   -------    0  -------   set
10   ----      ----    ----     0000.4000.4001 4.4.4.4                        Static Remote   -------    0 -------    set

Total number of entries are : 8

VTEP3#show nvo vxlan mac-table
=============================================================================================================================================
                                                VXLAN MAC Entries
=============================================================================================================================================
VNID Interface VlanId In-VlanId Mac-Addr       VTEP-Ip/ESI                  Type  Status  MAC move AccessPortDesc LeafFlag
---------------------------------------------------------------------------------------------------------------------------------------------

10  ----      ----    ----      0000.1000.1000 00:00:00:00:00:11:11:00:00:00  Static Remote  -------  0  -------    ----
10  ----      ----    ----      0000.1000.1001 00:00:00:00:00:11:11:00:00:00  Static Remote   ------- 0  -------    ----
10  ----      ----    ----      0000.2000.2000 2.2.2.2                        Static Remote   ------- 0   -------   ----
10  ----      ----    ----      0000.2000.2001 2.2.2.2                        Static Remote   -------  0 -------     ----
10  po2       1000    ----      0000.3000.3000 00:00:00:00:00:22:22:00:00:00  Static Local    -------  0 -------     set
10  po2      1000     ----      0000.3000.3001 00:00:00:00:00:22:22:00:00:00  Static Local    -------  0  -------    set
10  ----      ----    ----      0000.4000.4000 4.4.4.4                        Static Remote   -------   0  -------   set
```

```
10  ----      ----    ----      0000.4000.4001 4.4.4.4                    Static Remote   -------   0   ------- set

Total number of entries are : 8
```

Use the `show nvo vxlan arp-cache` command to verify the Address Resolution Protocol (ARP) cache information on all VTEP nodes. This command displays entries that map IPv4 addresses to MAC addresses within the specified VXLAN VNID network.

```
VTEP1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       100.100.100.1    0000.1000.1000 Static      Local      ----
10       103.103.103.1    0000.3000.3000 Static      Remote     ----
10       104.104.104.1    0000.4000.4000 Static    Remote     ----
10       200.200.200.1    0000.2000.2000 Static    Remote     ----
Total number of entries are 4

VTEP3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
10       100.100.100.1    0000.1000.1000 Static Remote  ----
10       103.103.103.1    0000.3000.3000 Static Local   ----
10       104.104.104.1    0000.4000.4000 Static Remote  ----
10       200.200.200.1    0000.2000.2000 Static Remote  ----
Total number of entries are 4
```

Use the `show nvo vxlan nd-cache` command to verify the Neighbor Discovery (ND) cache information on all VTEP nodes. This command displays entries that map IPv6 addresses to MAC addresses within the specified VXLAN VNID network.

```
VTEP1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type          Age-Out    Retries-Left
_____________________________________________________________________________
10       1000::1          0000.1000.1001 Static Local     ----
10       1003::1          0000.3000.3001 Static Remote    ----
10       1004::1          0000.4000.4001 Static Remote    ----
10       2000::1          0000.2000.2001 Static Remote    ----
Total number of entries are 4

VTEP3#show nvo vxlan nd-cache
VXLAN ND-CACHE Information
============================
VNID     Ip-Addr                                  Mac-Addr       Type          Age-Out
Retries-Left
_____________________________________________________________________________
_____________
10       1000::1                          0000.1000.1001 Static Remote    ----
10       1003::1                          0000.3000.3001 Static Local     ----
10       1004::1                          0000.4000.4001 Static Remote    ----
```

```
10       2000::1                          0000.2000.2001 Static Remote    ----
Total number of entries are 4
```

Network Topology Snippet Configurations

Here are the snippet configurations for all nodes in the given network topology.

VTEP1

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf VRF1
 rd 1.1.1.1:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po1
 switchport
 evpn multi-homed system-mac 0000.0000.1111
!
interface lo
 ip address 1.1.1.1/32 secondary
!
interface xe7
 channel-group 1 mode active
!
interface xe45
 ip address 10.10.10.1/24
!
interface xe49/2
 ip address 10.10.11.1/24
!
 exit
!

router ospf 100
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 10.10.10.0/24 area 0.0.0.0
 network 10.10.11.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 2.2.2.2 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 4.4.4.4 remote-as 100
```

```
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 2.2.2.2 activate
 neighbor 3.3.3.3 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan po1 1000
 map vnid 10
 mac 0000.1000.1000 ip 100.100.100.1
 mac 0000.1000.1001 ipv6 1000::1
!
```

VTEP2

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf VRF1
 rd 2.2.2.2:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po1
 switchport
 evpn multi-homed system-mac 0000.0000.1111
!
interface lo
 ip address 2.2.2.2/32 secondary
!
interface xe38
 channel-group 1 mode active
!
interface xe49/1
 ip address 20.20.20.1/24
!
interface xe50/1
 ip address 20.20.21.1/24
```

```
!
 exit
!

router ospf 100
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 20.20.20.0/24 area 0.0.0.0
 network 20.20.21.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 2.2.2.2
 neighbor 1.1.1.1 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 3.3.3.3 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan xe37 1000
 map vnid 10
 mac 0000.2000.2000 ip 200.200.200.1
 mac 0000.2000.2001 ipv6 2000::1
!
nvo vxlan access-if port-vlan po1 1000
 map vnid 10
 mac 0000.1000.1000 ip 100.100.100.1
 mac 0000.1000.1001 ipv6 1000::1
!
```

VTEP3

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn etree enable
!
mac vrf VRF1
 rd 3.3.3.3:100
 route-target both 100:100
```

```
!
nvo vxlan vtep-ip-global 3.3.3.3
!
nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po2
 switchport
 evpn multi-homed system-mac 0000.0000.2222
!
interface lo
 ip address 3.3.3.3/32 secondary
!
interface xe53/1
 ip address 30.30.30.1/24
!
interface xe54/1
 ip address 30.30.31.1/24
!
interface xe55/1
 channel-group 2 mode active
!
 exit
!
router ospf 100
 ospf router-id 3.3.3.3
 bfd all-interfaces
 network 3.3.3.3/32 area 0.0.0.0
 network 30.30.30.0/24 area 0.0.0.0
 network 30.30.31.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 3.3.3.3
 neighbor 1.1.1.1 remote-as 100
 neighbor 2.2.2.2 remote-as 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 4.4.4.4 update-source lo
 neighbor 4.4.4.4 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 2.2.2.2 activate
 neighbor 4.4.4.4 activate
 exit-address-family
 !
 exit
!
!
nvo vxlan access-if port-vlan po2 1000
 map vnid 10
 mac 0000.3000.3000 ip 103.103.103.1
```

```
 mac 0000.3000.3001 ipv6 1003::1
!
```

VTEP4

```
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn etree enable
!
mac vrf VRF1
 rd 4.4.4.4:100
 route-target both 100:100
!
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
 vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po2
 switchport
 evpn multi-homed system-mac 0000.0000.2222
!
interface lo
 ip address 4.4.4.4/32 secondary
!
interface xe11/1
 ip address 40.40.41.1/24
!
interface xe31/1
 channel-group 2 mode active
!
interface xe33
 ip address 40.40.40.1/24
!
interface xe34
 switchport
!
 exit
!
router ospf 100
 ospf router-id 4.4.4.4
 bfd all-interfaces
 network 4.4.4.4/32 area 0.0.0.0
 network 40.40.40.0/24 area 0.0.0.0
 network 40.40.41.0/24 area 0.0.0.0
!
router bgp 100
 bgp router-id 4.4.4.4
 neighbor 1.1.1.1 remote-as 100
```

```
 neighbor 2.2.2.2 remote-as 100
 neighbor 3.3.3.3 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 1.1.1.1 advertisement-interval 0
 neighbor 2.2.2.2 update-source lo
 neighbor 2.2.2.2 advertisement-interval 0
 neighbor 3.3.3.3 update-source lo
 neighbor 3.3.3.3 advertisement-interval 0
 !
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 2.2.2.2 activate
 neighbor 3.3.3.3 activate
 exit-address-family
 !
 exit
!
nvo vxlan access-if port-vlan xe34 1000
 map vnid 10
 mac 0000.4000.4000 ip 104.104.104.1
 mac 0000.4000.4001 ipv6 1004::1
!
nvo vxlan access-if port-vlan po2 1000
 map vnid 10
 mac 0000.3000.3000 ip 103.103.103.1
 mac 0000.3000.3001 ipv6 1003::1
!
```

SPINE1

```
!
qos enable
!
interface ce1/2
 ip address 40.40.40.2/24
!
interface ce1/4
 ip address 10.10.10.2/24
!
interface ce24/1
 ip address 30.30.30.2/24
!
interface ce27/1
 ip address 20.20.20.2/24
!
interface lo
 ip address 5.5.5.5/32 secondary
!
 exit
!
router ospf 100
 ospf router-id 5.5.5.5
 bfd all-interfaces
 network 5.5.5.5/32 area 0.0.0.0
 network 10.10.10.0/24 area 0.0.0.0
 network 20.20.20.0/24 area 0.0.0.0
```

```
 network 30.30.30.0/24 area 0.0.0.0
 network 40.40.40.0/24 area 0.0.0.0
!
```

SPINE2

```
!
qos enable
!
interface ce5/1
 ip address 20.20.21.2/24
!
interface ce10/1
 ip address 30.30.31.2/24
!
interface ce11/1
 ip address 40.40.41.2/24
!
interface ce14/2
 ip address 10.10.11.2/24
!
interface lo
 ip address 6.6.6.6/32 secondary
!
 exit
!
router ospf 100
 ospf router-id 6.6.6.6
 bfd all-interfaces
 network 6.6.6.6/32 area 0.0.0.0
 network 10.10.11.0/24 area 0.0.0.0
 network 20.20.21.0/24 area 0.0.0.0
 network 30.30.31.0/24 area 0.0.0.0
 network 40.40.41.0/24 area 0.0.0.0
!
```

SWITCH1

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
 vlan-reservation 4000-4094
 vlan 1000 bridge 1 state enable
!
interface po1
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe46
 channel-group 1 mode active
!
interface xe47
```

```
 channel-group 1 mode active
!
interface xe57
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
 exit
!
```

SWITCH2

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
 vlan-reservation 4000-4094
 vlan 1000 bridge 1 state enable
!
interface po2
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe33
 switchport
 bridge-group 1
 switchport mode hybrid
 switchport mode hybrid acceptable-frame-type all
 switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe49/1
 channel-group 2 mode active
!
interface xe51/1
 channel-group 2 mode active
!
 exit
!
```

Implementation Examples

Here is an example scenario and a solution for implementing EVPN E-Tree.

Scenario 1: Specific traffic isolation and control measures are essential in a network of EVPN L2VPN services or instances. Within a broadcast domain, services communicating with each other may result in flooding BUM traffic to all services within the domain. Moreover, hosts are learned and advertised between different sites/services.

Use Case 1: Implementing an EVPN E-Tree solution defines the network topology with distinct Root and Leaf classifications, BUM traffic flooding can be minimized, and traffic isolation can be achieved. This ensures efficient communication between services while preventing unnecessary traffic propagation and maintaining network integrity.

Scenario 2: An Internet Service Provider (ISP) provides services to multiple subscribers and aims to facilitate communication with them. However, the ISP needs to ensure that subscribers exclusively communicate with the ISP and not among themselves.

Use Case 2: Implementing EVPN E-Tree is essential to fulfill this requirement. By categorizing ISP services as Root and subscribers as Leaf, traffic isolation can be enforced. This configuration enables the ISP to communicate with subscribers while preventing inter-subscriber communication. As a result, network security is enhanced, and the ISP maintains control over communication within its network.

E-Tree CLI Commands

The EVPN E-Tree introduces the following configuration commands in OcNOS.

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration.

Command Syntax

```
evpn etree enable
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates how to activate E-Tree functionality for EVPN:

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable
```

Revised CLI Commands

The following is the revised command for configuring VXLAN EVPN E-Tree

nvo vxlan id

- The existing syntax now includes the newly added parameter for E-Tree, namely `etree-leaf`.
- The command `nvo vxlan id <VNID> ingress-replication inner-vid-disabled etree-leaf` allows users to tailor VXLAN behavior on a network device, specifying VXLAN parameters and indicating its

participation as a leaf node in an E-Tree deployment. For more details, refer to the nvo vxlan id command in the VXLAN Commands chapter in the *OcNOS VXLAN Guide*.

Troubleshooting

1. When traffic, whether unicast (UC) or broadcast, is passed to the Intra Leaf site:
   - Check the sub-interface or physical interface counters to monitor traffic throughput and potential issues.
   - Verify the Leaf status of the corresponding VNI to ensure proper functionality.
   - Use packet sniffing tools to analyze packets in the egress direction for any anomalies or errors.
   - MAC entries learned via leaf access port should include the `set` keyword in the MAC table output.
2. If UC traffic is routed within inter-PE leaf sites:
   - Check the Leaf status of the VNI at both participating PE devices to confirm operational status.
   - Check if the advertised MAC is in discard or non-discard status using the `show mac table` command and `l2 show` in the BCM shell.
3. Verify if BUM traffic is transmitted between Leaf sites inter-PE:
   - Ensure that a BUM tunnels are not established between inter-PE devices.
   - Validate this by examining the Multicast ingress group, using the `show evpn mpls tunnel` command. For EVPN MPLS, confirm that BUM tunnels are not created.
4. Investigate UC traffic drops from the Root to MH Leaf PE:
   - Check if MAC addresses are not installed in discard status within the MH peer's access port. This status could indicate issues with MAC learning or forwarding.
5. Evaluate traffic between Root and Leaf:
   - Confirm the establishment of both UC and BUM tunnels.
   - Ensure that unicast MAC addresses are not marked with a discard status in the MAC table.
6. Validate the exchange of routes between two BGP L2VPN peers:
   - Monitor BGP (Border Gateway Protocol) sessions to verify successful route exchange and propagation between the peers.
7. Convergence: Assess convergence by checking BFD configuration between BGP sessions.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Ethernet VPN Ethernet-Tree (EVPN E-Tree)	A networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, categorizing network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Virtual Extensible LAN (VXLAN)	A technology that provides encapsulation techniques to create virtualized Layer 2 networks over Layer 3 infrastructure, facilitating scalable and flexible network designs.
Ethernet Virtual Private Network (EVPN)	A Layer 2 VPN technology that extends Ethernet services across data centers and wide-area networks using BGP.
Multi-homing (MH)	The ability of a device to connect to multiple network segments simultaneously to increase network availability and redundancy.
Provider Edge (PE) Node	A device at the edge of a service provider network that connects to customer premises equipment (CE) and participates in providing services to customers.
Leaf Node	In the context of EVPN E-Tree, a network node categorized to handle communication within specific broadcast domains and may connect to Root nodes.
Root Node	A network node within EVPN E-Tree that serves as the central point of communication and handles BUM traffic distribution.
Ethernet Segment Identifier (ESI)	A unique identifier used to identify Ethernet segments within a VXLAN network.

CHAPTER 2 EVPN MPLS E-Tree

Overview

Ethernet VPN Ethernet-Tree (EVPN E-Tree), is a networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, especially in scenarios where specific services or devices need controlled communication. It categorizes network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Feature Characteristics

Implemented Scenario 1 of the EVPN E-Tree solution, as defined by RFC-8317, designates each Provider Edge (PE) node as either a Leaf or a Root site per Virtual Private Network (VPN) for VXLAN and MPLS EVPN in OcNOS.

Scenario 1: Leaf or Root Site(s) per PE

Scenario 1 involves a topology with three PE nodes: PE-1, PE-2, and PE-3. PE-1 and PE-2 are Multi-Homed nodes (MH-1 and MH-2), with PE-3 acting as the Root node. PE-1 and PE-2 function as Leaf nodes and are part of a single home access interface (SH1 and SH2).

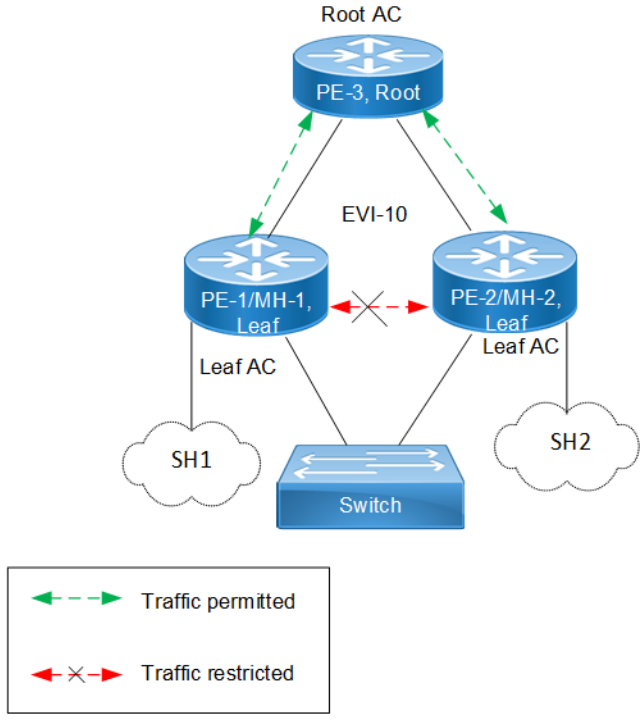

EVPN E-Tree

The classification ensures that communication follows specific rules:

- Communication between Leaf hosts is restricted, as indicated by red dotted lines with a cross mark (X) in the topology diagram. However, communication between Leaf and Root nodes, as well as between Root nodes, is permitted, marked by green dotted lines.

- Leaf nodes within PE-1 and PE-2 are isolated from each other, preventing intra-PE communication.

The scenario 1 is achieved through two main concepts:

1. **Inter-PE Communication**
   - The inter-PE Route Target (RT) Constraint Method is applicable only to Single-Homing (SH) devices. Two RTs per broadcast domain are utilized, with Leaf PEs exporting Leaf RTs and Root nodes exporting Root RTs. Leaf nodes import only Root RTs, allowing communication with Root PEs while preventing communication with other Leaf nodes. RT constraints limit the import of specific EVPN routes (MAC-IP and IMET routes) to designated paths for inter-PE communication.
   - IPI employs a proprietary method to support inter-PE connectivity for both SH and MH devices, using BGP extended community to advertise Leaf Indication in BGP routes and influence traffic flow for both Unicast and BUM traffic. This method enables implementation of ARP or ND cache suppression and MAC mobility sub-features specified in RFC-7432.
2. **Intra-PE communication:** Local Split Horizon controls intra-PE communication between Attachment Circuits (ACs) within Leaf PE nodes, ensuring that traffic between ACs does not egress to other Leaf ACs.

   Note: This functionality depends on hardware capabilities.

Benefits

EVPN E-Tree offers benefits in networking environments by providing efficient traffic control, enhanced security, scalability, and improved performance.

Efficient Traffic Control: EVPN E-Tree allows for efficient control over traffic within network broadcast domains. By segregating nodes into Leaf and Root categories, it enables precise management of communication flows, ensuring the traffic is directed only where needed.

Enhanced Security: The isolation of Leaf hosts from each other adds a layer of security to the network. This prevents unauthorized communication between devices within the same broadcast domain, reducing the risk of data breaches and unauthorized access.

Scalability: EVPN E-Tree is scalable, making it suitable for networks of various sizes and complexities. Whether deploying in small-scale environments or large enterprise networks, EVPN E-Tree offers flexibility and scalability to meet evolving business needs.

Improved Performance: By controlling communication paths and optimizing traffic flows, EVPN E-Tree can improve network performance. This ensures that critical data packets are delivered efficiently, reducing latency and enhancing overall network performance.

Prerequisites

In setting up a MPLS EVPN network, certain prerequisites are essential to ensure proper functionality and connectivity.

Ensure MPLS EVPN Configuration: Confirm that MPLS EVPN and MPLS MH filtering are already enabled in all leaf and root nodes of the network as they are required for MPLS EVPN Multihoming.

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn mpls multihoming enable
!
qos enable
```

```
!
```

Define Interfaces and Loopback Addresses: Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific system MAC addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!
interface po1
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.4321.1234
!
interface lo
 ip address 8.8.8.8/32 secondary
 ip router isis ISIS-IGP
 enable-ldp ipv4
!
interface xe8
 switchport
!
interface xe26
 channel-group 1 mode active
!
```

Configure ISIS and BGP for Dynamic Routing: Enable ISIS to facilitate dynamic routing on all Leaf and Root nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Additionally, establish BGP sessions to advertise routes between different nodes. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

```
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 8.8.8.8
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0008.00
!
router bgp 65535
 neighbor 9.9.9.9 remote-as 65535
 neighbor 24.24.24.24 remote-as 65535
 neighbor 26.26.26.26 remote-as 65535
 neighbor 29.29.29.29 remote-as 65535
 neighbor 9.9.9.9 update-source lo
 neighbor 9.9.9.9 fall-over bfd
 neighbor 24.24.24.24 update-source lo
 neighbor 24.24.24.24 fall-over bfd
 neighbor 26.26.26.26 update-source lo
 neighbor 26.26.26.26 fall-over bfd
```

```
 neighbor 29.29.29.29 update-source lo
 neighbor 29.29.29.29 fall-over bfd
 !
 address-family l2vpn evpn
 neighbor 9.9.9.9 activate
 neighbor 24.24.24.24 activate
 neighbor 26.26.26.26 activate
 neighbor 29.29.29.29 activate
 exit-address-family
 !
 exit
!
```

Configure LDP and RSVP for Efficient Network Operation: Enable Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) on all Leaf and Root nodes to optimize traffic routing and quality of service. LDP assigns labels for packet forwarding, while RSVP reserves network resources along specified paths to enhance network performance and reliability.

```
!
router ldp
 router-id 8.8.8.8
 fast-reroute
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 targeted-peer ipv4 9.9.9.9
  exit-targeted-peer-mode
 targeted-peer ipv4 24.24.24.24
  exit-targeted-peer-mode
 transport-address ipv4 8.8.8.8
!
router rsvp
!
rsvp-path LEAF1-ROOT2 mpls
 24.1.4.24 strict
!
rsvp-path LEAF1-ROOT1 mpls
 26.1.2.26 strict
!
rsvp-trunk LEAF1-ROOT1 ipv4
 primary fast-reroute protection facility
 primary path LEAF1-ROOT1
 to 9.9.9.9
!
rsvp-trunk LEAF1-ROOT2 ipv4
 primary fast-reroute protection facility
 primary path LEAF1-ROOT2
 to 24.24.24.24
!
```

Create VRF for Isolated Routing Instances: Configure VRF on all Leaf and Root nodes to create isolated routing instances within the network. This enables separate routing tables and forwarding behaviors for different groups of network resources.

```
!
mac vrf vrf103
 rd 8.8.8.8:103
 route-target both 65535:103
```

```
!
```

Connect Network Interfaces: Configure network interfaces on all Leaf and Root nodes with connection details, IP addresses, and protocol settings. Enable label-switching and configure participation in the ISIS routing protocol, including support for protocols like LDP and RSVP for IPv4. These configurations optimize routing and resource management across the network.

```
!
interface xe11
 description connected to ROOT2 int xe9
 ip address 24.1.4.25/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe20
 description connected to ROOT1 int xe20
 ip address 26.1.2.27/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
```

Configure Switch: Set up a VLAN bridge by enabling the VLAN and associating specific VLANs with the bridge. Configure network interfaces as trunk ports to allow traffic for all permitted VLANs across the network. Designate interfaces connected to Leaf and Root nodes as member ports of the VLAN bridge. This setup optimizes network segmentation and traffic management

```
!
bridge 1 protocol rstp vlan-bridge
!
vlan database
 vlan-reservation 4030-4094
 vlan 2-3010 bridge 1 state enable
!
interface po100
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
interface lo
 ip address 32.32.32.32/32 secondary
!
interface xe9
 channel-group 100 mode active
!
interface xe17
 channel-group 100 mode active
!
interface xe1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
 exit
```

!

Configuration

Configure various nodes within the topology to set up an MPLS EVPN E-Tree network, ensuring EVPN E-Tree for All-Active and Active-Standby redundancy and load balancing.

Topology

In the sample topology, Leaf nodes (LEAF1, LEAF2, LEAF3, and LEAF4), Root nodes (ROOT1 and ROOT2), and Switches (CE SWITCH, SWITCH1, and SWITCH2) form the network architecture. LEAF1 and LEAF2 are part of a Multi-homed group, with both connected to `po1` (MH2). LEAF1 and LEAF3 have single home access-if ports (SH3 and SH4, respectively). Similarly, ROOT1 and ROOT2 are part of a Multi-homed group with `po100` (MH1), and they each have a single home access-if port (SH1 and SH2, respectively). Leaf nodes are interconnected, and CE SWITCH, SWITCH1, and SWITCH2 are configured for Multi-homed connections to Leaf and Root nodes. SWITCH1 connects to LEAF1 and LEAF2, while CE SWITCH links to ROOT1 and ROOT2.

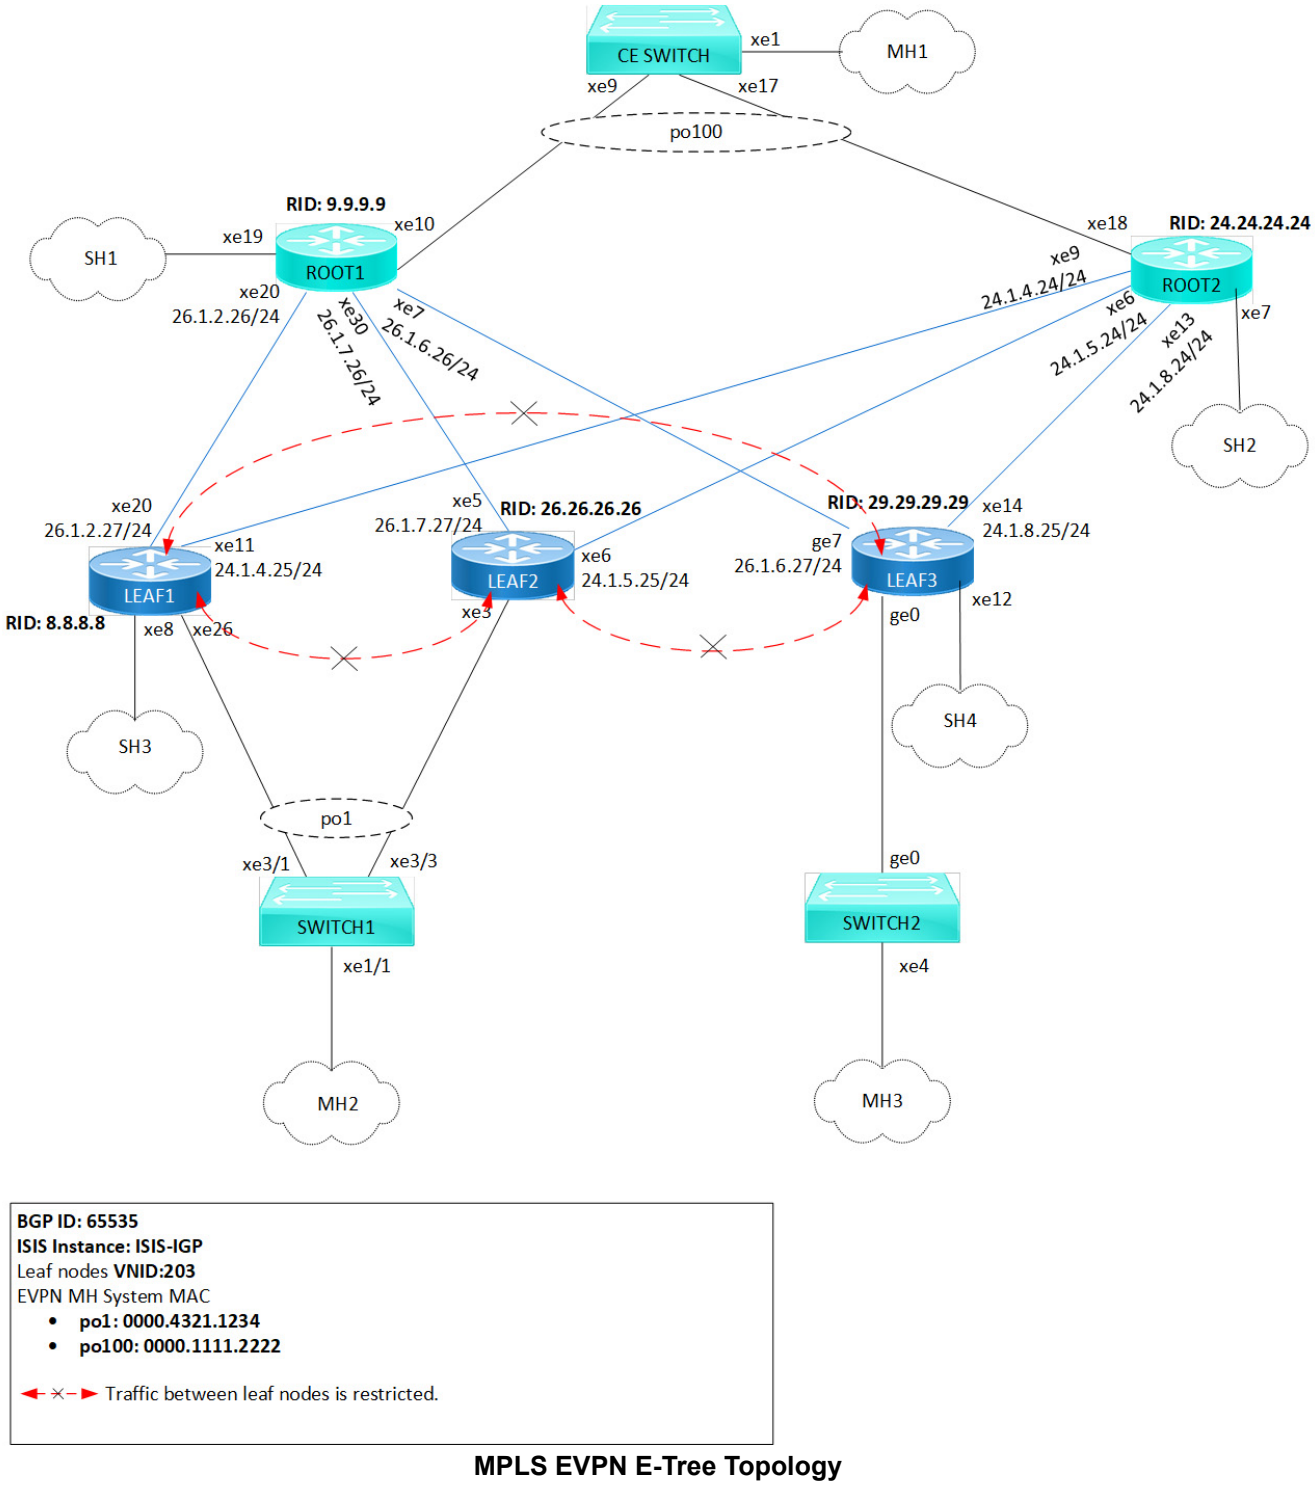

MPLS EVPN E-Tree Topology

Note: Before configuring E-Tree, meet all Prerequisites for the following nodes:

- Leaf nodes: LEAF1, LEAF2, and LEAF3
- Root nodes: ROOT1 and ROOT2
- Switches: CE SWITCH, SWITCH1 and SWITCH2

Enable EVPN E-Tree

The following E-Tree configurations applies to Leaf and Root nodes within the MPLS network.

1. Enable EVPN E-Tree which allows the nodes to participate in E-Tree functionality within the network, controlling traffic and establishing hierarchical connections between Leaf nodes in the network architecture.

   ```
   (config)#evpn etree enable
   ```

2. Set the MAC ageing time (`60 seconds`) to allow MAC addresses learned over EVPN MPLS to remain in the MAC table before timing out. Configure the global VTEP IP address (`8.8.8.8`) which serves as the global identifier for MPLS encapsulation and decapsulation within the network, facilitating proper communication and tunnel establishment.

   ```
   (config)#evpn mpls mac-ageing-time 60
   (config)#evpn mpls vtep-ip-global 8.8.8.8
   ```

3. Define MPLS identifier (`203`) to support hierarchical connectivity and traffic control within the EVPN MPLS network. On the EVPN MPLS node, specify EVPN-BGP as the host reachability protocol for the specified VRF (`vrf103`) to communicate and exchange reachability information within the network. To enable EVPN E-Tree on Leaf nodes, configure `etree-leaf` along with the MPLS identifier. This allows for efficient replication of traffic at the ingress point, optimizing the functionality of E-Tree Leaf nodes within the network architecture.

   ```
   (config)#evpn mpls id 203 etree-leaf
   (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf103
   (config-evpn-mpls)#exit
   ```

4. Enable port-VLAN mapping (`po1`) with VLAN ID (`103`) to facilitate multi-homed access. Enable EVPN functionality on the interface, allowing it to participate in MAC address distribution across the network.

   ```
   (config)#interface po1.103 switchport
   (config-if)#encapsulation dot1q 103
   (config-if)#load-interval 30
   (config-access-if)#access-if-evpn
   (config-access-if)#exit
   ```

Validation

Use the show commands described in this section to verify the network for proper MPLS EVPN E-Tree configuration.

Verify LDP sessions on all leaf and root nodes by using the `show ldp session` command. The `state` field (`OPERATIONAL`) indicates that the LDP session between the device and its peers is currently active.

```
LEAF1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       24.24.24.24        xe11       Passive   OPERATIONAL   30    01:13:29
       9.9.9.9            xe20       Passive   OPERATIONAL   30    01:13:29
```

Verify RSVP sessions on all leaf and root nodes by using the `show rsvp session` command. The `State` field (`UP`) indicates that the RSVP session between the ingress and egress routers is active and operational. Identify the different paths established within the network using the `LSPName` field.

```
LEAF1#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
```

```
Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                          State Uptime    Rt  Style  Labelin
Labelout
9.9.9.9         8.8.8.8         5001     2201     PRI    LEAF1-ROOT1-Primary            UP    01:13:16  1 1 SE     -       25601
24.24.24.24     8.8.8.8         5002     2202     PRI    LEAF1-ROOT2-Primary            UP    01:13:05  1 1 SE     -       25601
Total 2 displayed, Up 2, Down 0.

Egress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                          State Uptime    Rt  Style  Labelin
Labelout
8.8.8.8         9.9.9.9         5001     2201     PRI    ROOT1-LEAF1-Primary             UP    01:13:45  1 1 SE     25600    -
8.8.8.8         24.24.24.24     5001     2201     PRI    ROOT2-LEAF1-Primary             UP    01:13:24  1 1 SE     25601    -
Total 2 displayed, Up 2, Down 0.
```

Verify the BGP session status on all leaf and root nodes, using the `show bgp l2vpn evpn summary` command output. The `Up/Down` field indicates the duration for which the BGP session has been up or down.

```
LEAF1#show bgp l2vpn evpn summary
BGP router identifier 8.8.8.8, local AS number 65535
BGP table version is 33
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V AS     MsgRcv  MsgSen  TblVer  InQ   OutQ   Up/Down  State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
9.9.9.9       4 65535  514     443     33      0     0      01:13:53           114     59     5     50      0          0
24.24.24.24   4 65535  504     443     33      0     0      01:13:54           109     59     0     50      0          0
26.26.26.26   4 65535  322     391     33      0     0      01:13:23            49      0     0     49      0          0
29.29.29.29   4 65535  197     392     33      0     0      01:13:54             6      0     0      6      0          0

Total number of neighbors 4

Total number of Established sessions 4
```

Verify ESI information and the forwarding tunnel status on all leaf and root nodes, by examining the `show evpn mpls` command output. The DF- Status field displays the forwarding status as either a `Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF)`, and the ESI field displays the Ethernet Segment Identifier associated with each entry.

```
LEAF1#show evpn mpls
EVPN-MPLS Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VPN-ID   EVI-Name     EVI-Type Type Interface ESI                             VLAN      DF-Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________________________________
203      ----         L2       NW   ----      ----                          ----      ----      8.8.8.8          29.29.29.29
203      ----         L2       NW   ----      ----                          ----      ----      8.8.8.8          9.9.9.9
203      ----         L2       NW   ----      ----                          ----      ----      8.8.8.8          24.24.24.24
203      ----         L2       NW   ----      ----                          ----      ----      8.8.8.8          26.26.26.26
203      ----         --       AC   po1.103   00:00:00:43:21:12:34:00:00:00 ----      DF        ----             ----
203      ----         --       AC   po2.103   00:00:00:33:33:44:44:00:00:00 ----      DF        ----             ----

Total number of entries are 252
```

Static MAC-IP Advertisement

Configure static MAC-IP advertisement through SH and MH from Root and Leaf nodes. Advertise static MAC addresses for both IPv4 and IPv6 from all MH and SH nodes. Ensure that nodes within the same MH have identical MAC addresses configured under the port-channel access port.

Configure MH Nodes

Configure static MAC addresses for IPv4 (`30.30.30.3`) and IPv6 (`3000::1`) under the MH access-port (`po1`) with VLAN ID (`103`). Repeat the same configurations for other MH nodes using different static MAC addresses for both IPv4 and IPv6.

```
!
interface po1.103 switchport
 access-if-evpn
  map vpn-id 203
  mac 0000.7777.9999
  mac 0000.7777.6666 ip 30.30.30.3
  mac 0000.7777.6666 ipv6 3001::1
!
```

Configure SH Nodes

Configure static MAC addresses for IPv4 (`40.40.40.4`) and IPv6 (`4000::1`) under the SH access-port (`xe27`) with VLAN ID (`103`). This setup ensures that SH advertises these static MAC addresses over the specified access-port. Repeat the same configurations for other SH nodes using different static MAC addresses for both IPv4 and IPv6.

```
!
interface xe27.103 switchport
 encapsulation dot1q 100
 load-interval 30
 access-if-evpn
  map vpn-id 203
  mac 0000.0000.0011
  mac 0000.5544.4455 ip 40.40.40.4
  mac 0000.5544.4455 ipv6 4000::1
!
```

Validation

Verify the MAC table entries on MH nodes (MH1, MH2 and MH3) and the SH nodes (SH1, SH2, SH3, and SH4). MH nodes advertise their MAC addresses using the ESI values. Additionally, verify the IP addresses associated with SH nodes for MAC advertisement.

In the `show evpn mpls mac-table` command output, the MAC entries originated from Leaf Nodes will have the `LeafFlag` field status `set`.

Note:

- MAC IPv4 or IPv6 configured under SH Leaf node access port will be advertised to the Root nodes and other Leaf nodes.
- MAC IPv4 or IPv6 configured under an MH Leaf node access port must be symmetric and will be advertised to both the Root nodes and other leaf nodes.
- MAC IPv4 or IPv6 configured under either SH or MH Root node will be advertised to both the Root nodes and the Leaf nodes.
- The Leaf-to-Leaf communication will display MAC status and tunnel status per VNI as Leaf type. The MAC will be in the discard state in the BCM shell.

```
LEAF1#show evpn mpls mac-table
=============================================================================================================================
                                                EVPN MPLS MAC Entries
=============================================================================================================================
VNID Interface  VlanId  In-VlanId Mac-Addr       VTEP-Ip/ESI                    Type Status MAC move AccessPortDesc LeafFlag
-----------------------------------------------------------------------------------------------------------------------------

203  po1.103   ----   ----        0000.7777.9999 00:00:00:43:21:12:34:00:00:00      Static Local   ------- 0 -------  set
203  po1.103   ----   ----        0000.7777.6666 00:00:00:43:21:12:34:00:00:00      Static Local  -------  0 -------  set

Total number of entries are : 8

ROOT1#show evpn mpls mac-table
=============================================================================================================================
                                                EVPN MPLS MAC Entries
```

```
=============================================================================================================================
VNID Interface VlanId  In-VlanId Mac-Addr       VTEP-Ip/ESI                 Type Status  MAC move AccessPortDesc LeafFlag
-----------------------------------------------------------------------------------------------------------------------------

203 ----      ----     ----      0000.7777.9999 00:00:00:43:21:12:34:00:00:00  Static Remote -------    0 -------  set
203 ----      ----     ----      0000.7777.6666 00:00:00:43:21:12:34:00:00:00  Static Remote  -------  0  -------  set

Total number of entries are : 8
```

Use the `show evpn mpls arp-cache` command to verify the Address Resolution Protocol (ARP) cache information on all nodes. This command displays entries that map IPv4 addresses to MAC addresses within the specified EVPN ID network.

```
LEAF1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
================================
EVPN-ID   Ip-Addr          Mac-Addr        Type         Age-Out   Retries-Left
-----------------------------------------------------------------------------
203      30.30.30.3       0000.7777.6666 Static Local   ----
Total number of entries are 5

ROOT1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
================================
ARP Timeout :  570 sec   Random-Jitter-Max : 200

EVPN-ID   Ip-Addr          Mac-Addr        Type         Age-Out   Retries-Left
-----------------------------------------------------------------------------
203      30.30.30.3       0000.7777.6666 Static Remote  ----
Total number of entries are 5
```

Use the `show evpn mpls nd-cache` command to verify the Neighbor Discovery (ND) cache information on all nodes. This command displays entries that map IPv6 addresses to MAC addresses within the specified EVPN ID network.

```
LEAF1#show evpn mpls nd-cache
MPLS-EVPN ND-CACHE Information
================================
EVPN-ID  Ip-Addr                        Mac-Addr         Type         Age-Out Retries-Left
-----------------------------------------------------------------------------------------
203      3001::1                          0000.7777.6666 Static Local     ----
Total number of entries are 4

ROOT1#show evpn mpls nd-cache
MPLS-EVPN ND-CACHE Information
================================
EVPN-ID  Ip-Addr                        Mac-Addr         Type         Age-Out Retries-Left
-----------------------------------------------------------------------------------------
203      3001::1                          0000.7777.6666 Static Remote    ----
Total number of entries are 4
```

Network Topology Snippet Configurations

Here are the snippet configurations for all nodes in the given network topology.

LEAF1

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn etree enable
!
evpn mpls multihoming enable
!
mac vrf vrf103
 rd 8.8.8.8:103
 route-target both 65535:103
!
evpn mpls vtep-ip-global 8.8.8.8
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203 etree-leaf
 host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
router ldp
 router-id 8.8.8.8
 fast-reroute
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 targeted-peer ipv4 9.9.9.9
  exit-targeted-peer-mode
 targeted-peer ipv4 24.24.24.24
  exit-targeted-peer-mode
 transport-address ipv4 8.8.8.8
!
router rsvp
!
interface po1
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.4321.1234
!
interface po1.103 switchport
 encapsulation dot1q 103
 load-interval 30
 access-if-evpn
  map vpn-id 203
  mac 0000.7777.9999
  mac 0000.7777.6666 ip 30.30.30.3
  mac 0000.7777.6666 ipv6 3001::1
!
interface lo
 ip address 8.8.8.8/32 secondary
```

```
 ip router isis ISIS-IGP
 enable-ldp ipv4
!
interface xe8
 switchport
!
interface xe11
 description connected to ROOT2 int xe9
 ip address 24.1.4.25/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe20
 description connected to ROOT1 int xe20
 ip address 26.1.2.27/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe26
 channel-group 1 mode active
!
interface xe27
 speed 10g
!
interface xe27.100 switchport
 encapsulation dot1q 100
 load-interval 30
 access-if-evpn
  map vpn-id 200
  mac 0000.0000.0011
  mac 0000.5544.4455 ip 40.40.40.4
  mac 0000.5544.4455 ipv6 4000::1
!
 exit
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 8.8.8.8
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0008.00
!
router bgp 65535
 neighbor 9.9.9.9 remote-as 65535
```

```
 neighbor 24.24.24.24 remote-as 65535
 neighbor 26.26.26.26 remote-as 65535
 neighbor 29.29.29.29 remote-as 65535
 neighbor 9.9.9.9 update-source lo
 neighbor 9.9.9.9 fall-over bfd
 neighbor 24.24.24.24 update-source lo
 neighbor 24.24.24.24 fall-over bfd
 neighbor 26.26.26.26 update-source lo
 neighbor 26.26.26.26 fall-over bfd
 neighbor 29.29.29.29 update-source lo
 neighbor 29.29.29.29 fall-over bfd
 !
 address-family l2vpn evpn
 neighbor 9.9.9.9 activate
 neighbor 24.24.24.24 activate
 neighbor 26.26.26.26 activate
 neighbor 29.29.29.29 activate
 exit-address-family
 !
 exit
!
rsvp-path LEAF1-ROOT2 mpls
 24.1.4.24 strict
!
rsvp-path LEAF1-ROOT1 mpls
 26.1.2.26 strict
!
rsvp-trunk LEAF1-ROOT1 ipv4
 primary fast-reroute protection facility
 primary path LEAF1-ROOT1
 to 9.9.9.9
!
rsvp-trunk LEAF1-ROOT2 ipv4
 primary fast-reroute protection facility
 primary path LEAF1-ROOT2
 to 24.24.24.24
!
```

LEAF2

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn mpls multihoming enable
!
mac vrf vrf103
 rd 26.26.26.26:103
 route-target both 65535:103
!
evpn mpls vtep-ip-global 26.26.26.26
!
evpn mpls mac-ageing-time 60
```

```
!
evpn mpls id 203 etree-leaf
 host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
router ldp
 router-id 26.26.26.26
 fast-reroute
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 targeted-peer ipv4 9.9.9.9
  exit-targeted-peer-mode
 targeted-peer ipv4 24.24.24.24
  exit-targeted-peer-mode
 transport-address ipv4 26.26.26.26
!
router rsvp
!
interface po1
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.4321.1234
!
interface po1.103 switchport
 encapsulation dot1q 103
 load-interval 30
 access-if-evpn
  map vpn-id 203
!
interface lo
 ip address 26.26.26.26/32 secondary
 ip router isis ISIS-IGP
 enable-ldp ipv4
!
interface xe3
 channel-group 1 mode active
!
interface xe5
 description connected to ROOT1 int xe30
 ip address 26.1.7.27/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe6
 description connected to ROOT2 int xe6
 ip address 24.1.5.25/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
 exit
```

```
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 26.26.26.26
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0026.00
!
router bgp 65535
 neighbor 8.8.8.8 remote-as 65535
 neighbor 9.9.9.9 remote-as 65535
 neighbor 24.24.24.24 remote-as 65535
 neighbor 29.29.29.29 remote-as 65535
 neighbor 8.8.8.8 update-source lo
 neighbor 8.8.8.8 fall-over bfd
 neighbor 9.9.9.9 update-source lo
 neighbor 9.9.9.9 fall-over bfd
 neighbor 24.24.24.24 update-source lo
 neighbor 24.24.24.24 fall-over bfd
 neighbor 29.29.29.29 update-source lo
 neighbor 29.29.29.29 fall-over bfd
 !
 address-family l2vpn evpn
 neighbor 8.8.8.8 activate
 neighbor 9.9.9.9 activate
 neighbor 24.24.24.24 activate
 neighbor 29.29.29.29 activate
 exit-address-family
 !
 exit
!
rsvp-path LEAF2-ROOT2 mpls
 24.1.5.24 strict
!
rsvp-path LEAF2-ROOT1 mpls
 26.1.7.26 strict
!
rsvp-trunk LEAF2-ROOT1 ipv4
 primary fast-reroute protection facility
 primary path LEAF2-ROOT1
 to 9.9.9.9
!
rsvp-trunk LEAF2-ROOT2 ipv4
 primary fast-reroute protection facility
 primary path LEAF2-ROOT2
 to 24.24.24.24
!
```

LEAF3

```
!
evpn mpls enable
!
mac vrf vrf103
 rd 29.29.29.29:103
 route-target both 65535:103
!
evpn mpls vtep-ip-global 29.29.29.29
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203 etree-leaf
 host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
router ldp
 router-id 29.29.29.29
 fast-reroute
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 targeted-peer ipv4 9.9.9.9
  exit-targeted-peer-mode
 targeted-peer ipv4 24.24.24.24
  exit-targeted-peer-mode
 transport-address ipv4 29.29.29.29
!
router rsvp
!
interface ge0
 static-channel-group 3
!
interface ge7
 description connected to ROOT1 int xe7
 ip address 26.1.6.27/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface lo
 ip address 29.29.29.29/32 secondary
 ip router isis ISIS-IGP
 enable-ldp ipv4
!
interface xe12
 switchport
!
interface xe12.103 switchport
 encapsulation dot1q 103
 load-interval 30
 access-if-evpn
  map vpn-id 203
```

```
!
interface xe14
 description connected to ROOT2 int xe13
 ip address 24.1.8.25/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
 exit
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 29.29.29.29
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0029.00
!
router bgp 65535
 neighbor 8.8.8.8 remote-as 65535
 neighbor 9.9.9.9 remote-as 65535
 neighbor 24.24.24.24 remote-as 65535
 neighbor 26.26.26.26 remote-as 65535
 neighbor 8.8.8.8 update-source lo
 neighbor 8.8.8.8 fall-over bfd
 neighbor 9.9.9.9 update-source lo
 neighbor 9.9.9.9 fall-over bfd
 neighbor 24.24.24.24 update-source lo
 neighbor 24.24.24.24 fall-over bfd
 neighbor 26.26.26.26 update-source lo
 neighbor 26.26.26.26 fall-over bfd
 !
 address-family l2vpn evpn
 neighbor 8.8.8.8 activate
 neighbor 9.9.9.9 activate
 neighbor 24.24.24.24 activate
 neighbor 26.26.26.26 activate
 exit-address-family
 !
 exit
!
rsvp-path LEAF3-ROOT2 mpls
 24.1.8.24 strict
!
rsvp-path LEAF3-ROOT1 mpls
 26.1.6.26 strict
!
rsvp-trunk LEAF3-ROOT1 ipv4
```

```
 primary fast-reroute protection facility
 primary path LEAF3-ROOT1
 to 9.9.9.9
!
rsvp-trunk LEAF3-ROOT2 ipv4
 primary fast-reroute protection facility
 primary path LEAF3-ROOT2
 to 24.24.24.24
!
```

ROOT1

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn mpls multihoming enable
!
mac vrf vrf103
 rd 9.9.9.9:103
 route-target both 65535:103
!
evpn mpls vtep-ip-global 9.9.9.9
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203
 host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
bridge 1 protocol rstp vlan-bridge
!
router ldp
 router-id 9.9.9.9
 fast-reroute
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 targeted-peer ipv4 8.8.8.8
  exit-targeted-peer-mode
 targeted-peer ipv4 26.26.26.26
  exit-targeted-peer-mode
 transport-address ipv4 9.9.9.9
!
router rsvp
!
interface po100
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.1111.2222
!
interface po100.103 switchport
```

```
 encapsulation dot1q 103
 load-interval 30
 access-if-evpn
  map vpn-id 203
!
interface lo
 ip address 9.9.9.9/32 secondary
 ip router isis ISIS-IGP
 enable-ldp ipv4
!
interface xe7
 description connected to LEAF3 int ge7
 speed 1g
 ip address 26.1.6.26/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe10
 channel-group 100 mode active
!
interface xe17.100 switchport
 description for Static mac advertize
 encapsulation dot1q 100
 load-interval 30
 access-if-evpn
  map vpn-id 200
  mac 0000.0000.0022
  mac 0000.00dc.0001 ip 10.10.10.1
  mac 0000.00dc.0001 ipv6 1001::1
!
interface xe19
 switchport
!
interface xe20
 description connected to LEAF1 int xe20
 ip address 26.1.2.26/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe30
 description connected to LEAF2 int xe5
 speed 10g
 ip address 26.1.7.26/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
 exit
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
```

```
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 9.9.9.9
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0009.00
!
router bgp 65535
 neighbor 8.8.8.8 remote-as 65535
 neighbor 24.24.24.24 remote-as 65535
 neighbor 26.26.26.26 remote-as 65535
 neighbor 29.29.29.29 remote-as 65535
 neighbor 8.8.8.8 update-source lo
 neighbor 8.8.8.8 fall-over bfd
 neighbor 24.24.24.24 update-source lo
 neighbor 24.24.24.24 fall-over bfd
 neighbor 26.26.26.26 update-source lo
 neighbor 26.26.26.26 fall-over bfd
 neighbor 29.29.29.29 update-source lo
 neighbor 29.29.29.29 fall-over bfd
 !
 address-family l2vpn evpn
 neighbor 8.8.8.8 activate
 neighbor 24.24.24.24 activate
 neighbor 26.26.26.26 activate
 neighbor 29.29.29.29 activate
 exit-address-family
 !
 exit
!
rsvp-path ROOT1-LEAF3 mpls
 26.1.6.27 strict
!
rsvp-path ROOT1-LEAF2 mpls
 26.1.7.27 strict
!
rsvp-path ROOT1-LEAF1 mpls
 26.1.2.27 strict
!
rsvp-trunk ROOT1-LEAF1 ipv4
 primary fast-reroute protection facility
 primary path ROOT1-LEAF1
 to 8.8.8.8
!
rsvp-trunk ROOT1-LEAF2 ipv4
 primary fast-reroute protection facility
 primary path ROOT11-LEAF2
 to 26.26.26.26
!
rsvp-trunk ROOT1-LEAF3 ipv4
 primary fast-reroute protection facility
```

```
 primary path ROOT1-LEAF3
 to 29.29.29.29
!
```

ROOT2

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn mpls multihoming enable
!
mac vrf vrf103
 rd 24.24.24.24:103
 route-target both 65535:103
!
evpn mpls vtep-ip-global 24.24.24.24
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203
 host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
router ldp
 router-id 24.24.24.24
 fast-reroute
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 targeted-peer ipv4 8.8.8.8
  exit-targeted-peer-mode
 targeted-peer ipv4 26.26.26.26
  exit-targeted-peer-mode
 transport-address ipv4 24.24.24.24
!
router rsvp
!
interface po100
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.1111.2222
!
interface po100.103 switchport
 encapsulation dot1q 103
 load-interval 30
 access-if-evpn
  map vpn-id 203
!
interface lo
 ip address 24.24.24.24/32 secondary
 ip router isis ISIS-IGP
```

```
 enable-ldp ipv4
!
interface xe6
 description connected to LEAF2 int xe6
 speed 10g
 ip address 24.1.5.24/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe7
 switchport
!
interface xe9
 description connected to LEAF1 int xe11
 speed 10g
 ip address 24.1.4.24/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe13
 description connected to LEAF3 int xe14
 speed 10g
 ip address 24.1.8.24/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe18
 channel-group 100 mode active
!
 exit
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 24.24.24.24
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0024.00
!
router bgp 65535
 neighbor 8.8.8.8 remote-as 65535
 neighbor 9.9.9.9 remote-as 65535
 neighbor 26.26.26.26 remote-as 65535
```

```
 neighbor 29.29.29.29 remote-as 65535
 neighbor 8.8.8.8 update-source lo
 neighbor 8.8.8.8 fall-over bfd
 neighbor 9.9.9.9 update-source lo
 neighbor 9.9.9.9 fall-over bfd
 neighbor 26.26.26.26 update-source lo
 neighbor 26.26.26.26 fall-over bfd
 neighbor 29.29.29.29 update-source lo
 neighbor 29.29.29.29 fall-over bfd
 !
 address-family l2vpn evpn
 neighbor 8.8.8.8 activate
 neighbor 9.9.9.9 activate
 neighbor 26.26.26.26 activate
 neighbor 29.29.29.29 activate
 exit-address-family
 !
 exit
!
rsvp-path ROOT2-LEAF1 mpls
 24.1.4.25 strict
!
rsvp-path ROOT2-LEAF2 mpls
 24.1.5.25 strict
!
rsvp-path ROOT2-LEAF3 mpls
 24.1.8.25 strict
!
rsvp-trunk ROOT2-LEAF1 ipv4
 primary fast-reroute protection facility
 primary path ROOT2-LEAF1
 to 8.8.8.8
!
rsvp-trunk ROOT2-LEAF2 ipv4
 primary fast-reroute protection facility
 primary path ROOT2-LEAF2
 to 26.26.26.26
!
rsvp-trunk ROOT2-LEAF3 ipv4
 primary fast-reroute protection facility
 primary path ROOT2-LEAF3
 to 29.29.29.29
!
```

CE SWITCH

```
!
bridge 1 protocol rstp vlan-bridge
!
vlan database
 vlan-reservation 4030-4094
 vlan 2-3010 bridge 1 state enable
!
interface po100
 switchport
 bridge-group 1
 switchport mode trunk
```

```
 switchport trunk allowed vlan all
!
interface lo
 ip address 32.32.32.32/32 secondary
!
interface xe9
 channel-group 100 mode active
!
interface xe17
 channel-group 100 mode active
!
interface xe1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
 exit
!
```

SWITCH1

```
!
bridge 1 protocol rstp vlan-bridge
!
 vlan-reservation 4020-4062
 vlan 2-3000 bridge 1 state enable
!
interface po1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
interface lo
 ip address 7.7.7.7/32 secondary
!
interface xe1/1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
interface xe3/1
 channel-group 1 mode active
!
interface xe3/3
 channel-group 1 mode active
!
 exit
!
```

SWITCH2

```
!
bridge 1 protocol rstp vlan-bridge
!
```

```
vlan database
 vlan 2-3000 bridge 1 state enable
!
interface sa3
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
interface ge0
 static-channel-group 3
!
interface lo
 ip address 23.23.23.23/32 secondary
!
interface xe4
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
```

E-Tree Active-Standby Configuration

To set up an E-Tree network with Active-Standby redundancy and load balancing, follow these steps:

- Connect the Switch (P1) to the Root1, LEAF1, and LEAF2 nodes in the MPLS EVPN E-Tree Topology.
- Set up the VRF, EVPN, Port-Active, and Single-Active Redundancy configuration on Root MH and Leaf MH nodes.

For more details on Active-Standby configuration, refer to the section EVPN Active-Standby.

LEAF1

```
!
mac vrf vrf600
 rd 26.26.26.26:600
 route-target both 65535:600
!
evpn mpls id 681 etree-leaf
 host-reachability-protocol evpn-bgp vrf600
!
interface po1
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.4321.1234 load-balancing port-active
  service-carving auto
!
interface po1.681 switchport
 encapsulation dot1q 681
 load-interval 30
 access-if-evpn
  map vpn-id 681
!
interface sa1
 switchport
 load-interval 30
 evpn multi-homed esi 11:22:33:00:00:00:55:66:77 load-balancing single-active
```

```
  service-carving auto
!
interface sa1.681 switchport
 encapsulation dot1q 681
 load-interval 30
 access-if-evpn
  map vpn-id 681
!
interface xe4
 description connected to P1 int xe43
 speed 10g
 load-interval 30
 ip address 25.1.2.25/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
```

LEAF2

```
!
mac vrf vrf600
 rd 26.26.26.26:600
 route-target both 65535:600
!
evpn mpls id 681 etree-leaf
 host-reachability-protocol evpn-bgp vrf600
!
interface po1
 switchport
 load-interval 30
 evpn multi-homed system-mac 0000.4321.1234 load-balancing port-active
  service-carving auto
!
interface po1.681 switchport
 encapsulation dot1q 681
 load-interval 30
 access-if-evpn
  map vpn-id 681
!
interface sa2
 switchport
 load-interval 30
 evpn multi-homed esi 11:22:33:00:00:00:55:66:77 load-balancing single-active
  service-carving auto
!
interface sa2.681 switchport
 encapsulation dot1q 681
 load-interval 30
 access-if-evpn
  map vpn-id 681
!
interface xe21
 description connected to P1 int xe43
 speed 10g
 load-interval 30
```

```
 ip address 27.1.2.25/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
```

P1

```
!
router ldp
 router-id 6.6.6.6
 graceful-restart full
 graceful-restart timers neighbor-liveness 120
 graceful-restart timers max-recovery 120
 session-protection duration 10
 transport-address ipv4 6.6.6.6
!
interface lo
 ip address 127.0.0.1/8
 ip address 6.6.6.6/32 secondary
 ipv6 address ::1/128
 ip router isis ISIS-IGP
 enable-ldp ipv4
!
interface xe43
 description connected to LEAF1 int xe4
 speed 10g
 load-interval 30
 ip address 25.1.2.24/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe45
 description connected to ROOT1 int xe2
 speed 10g
 load-interval 30
 ip address 26.1.3.27/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
interface xe47
 description connected to LEAF2 int xe21
 speed 10g
 load-interval 30
 ip address 27.1.2.24/24
 label-switching
 ip router isis ISIS-IGP
 enable-ldp ipv4
 enable-rsvp
!
 exit
```

```
!
router isis ISIS-IGP
 is-type level-1
 authentication mode md5 level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 6.6.6.6
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0006.00
!
```

Validation

To verify the status of the ESI, whether it's active or standby, use the `show evpn load-balance all` command. This command helps debug and understand if the election process is occurring correctly. For the ESI `00:00:00:43:21:12:34:00:00:00`, LEAF1 is active, and LEAF2 is on standby in port-active mode. For the ESI `00:11:22:33:00:00:00:55:66:77`, LEAF2 is active, and LEAF1 is on standby in single-active mode.

```
LEAF1#show evpn load-balance all
ESI                             AC-IF/PE     PE-IP-ADDRESS    Redundancy   Service-carving weight Revertive  AC-DF   Status
=============================================================================================================================
00:00:00:43:21:12:34:00:00:00   LOCAL      8.8.8.8          port-active    auto            0      NO         NA     ACTIVE
00:00:00:43:21:12:34:00:00:00   REMOTE     26.26.26.26      port-active    auto            0      NO         NA     STANDBY
00:11:22:33:00:00:00:55:66:77   sa1.681    8.8.8.8          single-active  auto            0      NO         NO     STANDBY


LEAF2#show evpn load-balance all
ESI                             AC-IF/PE     PE-IP-ADDRESS    Redundancy   Service-carving weight Revertive  AC-DF   Status
=============================================================================================================================
00:00:00:43:21:12:34:00:00:00   REMOTE     8.8.8.8          port-active   auto            0      NO         NA     ACTIVE
00:00:00:43:21:12:34:00:00:00   LOCAL      26.26.26.26      port-active   auto            0      NO         NA     STANDBY
00:11:22:33:00:00:00:55:66:77   sa2.681    26.26.26.26      single-active  auto           0      NO         NO     ACTIVE
```

All MAC addresses in Root and Leaf nodes will be synchronized.

```
LEAF1#show evpn mpls mac-table
=============================================================================================================================
                                                EVPN MPLS MAC Entries
=============================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                 Type           Status
MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________________________________________________________

681        po1.681   ----      ----      0000.da00.0001 00:00:00:43:21:12:34:00:00:00                Dynamic Local   ------
-    0        -------          set
681        ----      ----      ----      0000.ea00.0001 00:00:00:11:11:22:22:00:00:00                Dynamic Remote  ------
-    0        -------          ----

Total number of entries are : 2

LEAF2#show evpn mpls mac-table
=============================================================================================================================
                                                EVPN MPLS MAC Entries
=============================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                 Type           Status
MAC move AccessPortDesc LeafFlag
_____________________________________________________________________________________________________________________________
```

```
681        ----      ----      ----      0000.da00.0001 00:00:00:43:21:12:34:00:00:00                Dynamic Remote  ------
-    0        -------          set
681        ----      ----      ----      0000.ea00.0001 00:00:00:11:11:22:22:00:00:00                Dynamic Remote  ------
-    0        -------          ----

Total number of entries are : 2

ROOT1#show evpn mpls mac-table
=============================================================================================================================
                                                    EVPN MPLS MAC Entries
=============================================================================================================================
VNID       Interface VlanId    In-VlanId Mac-Addr       VTEP-Ip/ESI                                 Type            Status
MAC move AccessPortDesc LeafFlag
-----------------------------------------------------------------------------------------------------------------------------

681        ----      ----      ----      0000.da00.0001 00:00:00:43:21:12:34:00:00:00                Dynamic Remote  ------
-    0        -------          set
681        po100.681 ----      ----      0000.ea00.0001 00:00:00:11:11:22:22:00:00:00                Dynamic Local   ------
-    0        -------          ----

Total number of entries are : 2
```

Implementation Examples

Here is an example scenario and a solution for implementing EVPN E-Tree.

Scenario 1: Specific traffic isolation and control measures are essential in a network of EVPN L2VPN services or instances. Within a broadcast domain, services communicating with each other may result in flooding BUM traffic to all services within the domain. Moreover, hosts are learned and advertised between different sites/services.

Use Case 1: Implementing an EVPN E-Tree solution defines the network topology with distinct Root and Leaf classifications, BUM traffic flooding can be minimized, and traffic isolation can be achieved. This ensures efficient communication between services while preventing unnecessary traffic propagation and maintaining network integrity.

Scenario 2: An Internet Service Provider (ISP) provides services to multiple subscribers and aims to facilitate communication with them. However, the ISP needs to ensure that subscribers exclusively communicate with the ISP and not among themselves.

Use Case 2: Implementing EVPN E-Tree is essential to fulfill this requirement. By categorizing ISP services as Root and subscribers as Leaf, traffic isolation can be enforced. This configuration enables the ISP to communicate with subscribers while preventing inter-subscriber communication. As a result, network security is enhanced, and the ISP maintains control over communication within its network.

E-Tree CLI Commands

The EVPN E-Tree introduces the following configuration commands in OcNOS.

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration.

Command Syntax

```
evpn etree enable
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates how to activate E-Tree functionality for EVPN:

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable
```

Revised CLI Commands

The following is the revised command for configuring MPLS EVPN E-Tree

evpn mpls id

- The existing syntax now includes the newly added parameter for E-Tree, namely `etree-leaf`.
- The command `evpn mpls id <ID> etree-leaf` allows users to tailor MPLS EVPN behavior on a network device, indicating its participation as a leaf node in an E-Tree deployment. For more details, refer to the evpn mpls id command in the EVPN MPLS Commands chapter in the *OcNOS Multi-Protocol Label Switching Guide*.

Troubleshooting

1. When traffic, whether unicast (UC) or broadcast, is passed to the Intra Leaf site:
   - Check the sub-interface or physical interface counters to monitor traffic throughput and potential issues.
   - Verify the Leaf status of the corresponding VNI to ensure proper functionality.
   - Use packet sniffing tools to analyze packets in the egress direction for any anomalies or errors.
   - MAC entries learned via leaf access port should include the `set` keyword in the MAC table output.
2. If UC traffic is routed within inter-PE leaf sites:
   - Check the Leaf status of the VNI at both participating PE devices to confirm operational status.
   - Check if the advertised MAC is in discard or non-discard status using the `show mac table` command and `l2 show` in the BCM shell.
3. Verify if BUM traffic is transmitted between Leaf sites inter-PE:
   - Ensure that a BUM tunnels are not established between inter-PE devices.
   - Validate this by examining the Multicast ingress group, using the `show evpn mpls tunnel` command. For EVPN MPLS, confirm that BUM tunnels are not created.
4. Investigate UC traffic drops from the Root to MH Leaf PE:

- Check if MAC addresses are not installed in discard status within the MH peer's access port. This status could indicate issues with MAC learning or forwarding.

5. Evaluate traffic between Root and Leaf:
   - Confirm the establishment of both UC and BUM tunnels.
   - Ensure that unicast MAC addresses are not marked with a discard status in the MAC table.
6. Validate the exchange of routes between two BGP L2VPN peers:
   - Monitor BGP (Border Gateway Protocol) sessions to verify successful route exchange and propagation between the peers.
7. Convergence: Assess convergence by checking BFD configuration between BGP sessions.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
EVPN E-Tree (Ethernet VPN Ethernet-Tree)	A networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, categorizing network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.
EVPN (Ethernet Virtual Private Network)	A Layer 2 VPN technology that extends Ethernet services across data centers and wide-area networks using BGP.
Multi-homing (MH)	The ability of a device to connect to multiple network segments simultaneously to increase network availability and redundancy.
Provider Edge (PE) Node	A device at the edge of a service provider network that connects to customer premises equipment (CE) and participates in providing services to customers.
Leaf Node	In the context of EVPN E-Tree, a network node categorized to handle communication within specific broadcast domains and may connect to Root nodes.
Root Node	A network node within EVPN E-Tree that serves as the central point of communication and handles BUM traffic distribution.
Ethernet Segment Identifier (ESI)	A unique identifier used to identify Ethernet segments within a MPLS network.

CHAPTER 3 LDP Tunneling over RSVP-TE

Overview

LDP-over-RSVP-TE tunneling is a technique used in MPLS networks to combine the strengths of Label Distribution Protocol (LDP) and Resource Reservation Protocol Traffic Engineering (RSVP-TE). This approach allows LDP Label Switched Paths (LSPs) to be encapsulated within RSVP-TE LSPs, providing enhanced traffic engineering capabilities while maintaining operational simplicity.

Feature Characteristics

LDP-over-RSVP-TE facilitates the integration of LDP LSPs within RSVP-TE tunnels, leveraging the strengths of both protocols. It harnesses RSVP-TE's traffic engineering capabilities for path computation, bandwidth reservation, and quality of service (QoS) provisioning. Ingress nodes execute FEC resolution to designate the suitable RSVP-TE tunnel for tunneling LDP LSPs, establishing hierarchical LSPs with RSVP-TE as the outer label and LDP as the inner label.

Benefits

LDP-over-RSVP-TE offers significant benefits are:

- **Advanced Traffic Engineering**: By leveraging RSVP-TE's advanced traffic engineering mechanisms, LDP-over-RSVP-TE enables efficient path computation, bandwidth reservation, and Quality of Service (QoS) provisioning.
- **Simplified Network Topology**: eliminates the need for a full mesh of intra-area RSVP LSPs (Label Switched Paths) between PE (Provider Edge) nodes.
- **Enhanced Resilience with Fast Reroute (FRR)**: Inherit RSVP-TE's Fast Reroute (FRR) capabilities. This means that in case of link or node failures, the network can quickly reroute traffic along pre-established backup paths
- **Flexible Hierarchical LSP Design**: Provides flexibility in network design by allowing for hierarchical LSPs (Label Switched Paths) where RSVP-TE serves as the outer label and LDP as the inner label.

Prerequisites

Before configuring this feature, ensure the following:

- A functional MPLS network with support for both LDP and RSVP-TE protocols.
- Network devices (routers or switches) capable of supporting LDP and RSVP-TE functionalities.

Limitations

The limitations are:

- LDP-over-RSVP tunneling is supported only with ISIS as IGP.
- Tunneling over inter-domain IGP area is not supported.
- LDP LSP tunneling over RSVP multipath is not supported.
- MPLS trace route is not supported in LDP-over-RSVP tunneling path.
- Dynamic TLDP sessions are not supported, TLDP session has to be explicitly configured.

- LFA and/or RLFA protection is not supported for LDP-over-RSVP tunnels.
- MPLS EVPN ELAN services over LoR are not supported.

Configuration for LDP Tunneling Over RSVP

Configure various nodes within the topology to set up a LDP Tunneling over RSVP session.

Topology

This sample topology provides basic connectivity and routing between the devices.

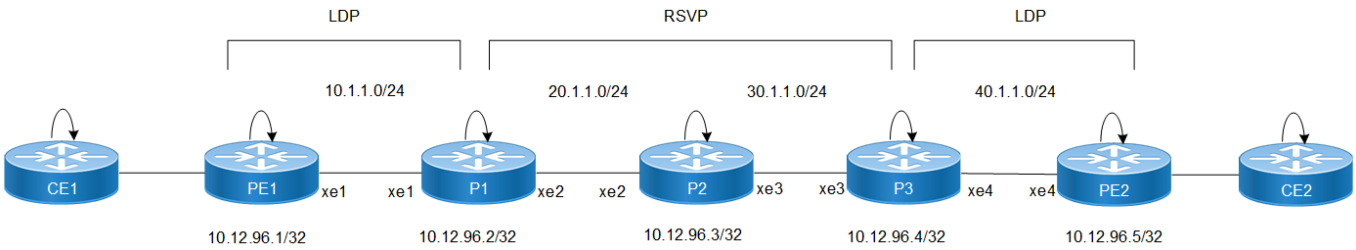

LDP Tunneling over RSVP Configuration

Configure LDP Tunneling over RSVP on PE1 Router

Follow the steps to configure the LDP tunneling over RSVP on PE1 router:

1. Configure the loopback interface with an IP address.

```
PE1(config)#interface lo
PE1(config-if)#ip address 10.12.96.1/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
PE1(config)#router ldp
PE1(config-router)#router-id 10.12.96.1
PE1(config-router)#transport-address ipv4 10.12.96.1
```

3. Configure global RSVP parameters.

```
PE1(config)#router rsvp
```

4. Configure the interface facing the network side with an IP address, enable label switching, and enable LDP.

```
PE1(config)#interface xe1
PE1(config-if)#ip address 10.1.1.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
```

5. If using ISIS as the Interior Gateway Protoco (IGP), configure ISIS parameters including traffic engineering.

```
PE1(config)#router isis ISIS-IGP
PE1(config-router)#is-type level-1
PE1(config-router)#metric-style wide
PE1(config-router)#mpls traffic-eng router-id 10.12.96.1
PE1(config-router)#mpls traffic-eng level-1
PE1(config-router)#capability cspf
PE1(config-router)#dynamic-hostname
PE1(config-router)#net 49.0000.0000.0001.00
```

```
PE1(config-router)#exit
```

Configure LDP Tunneling over RSVP on P1 Router

Follow the steps to configure the LDP tunneling over RSVP on P1 router:

1. Configure the loopback interface with an IP address.
```
P1(config)#interface lo
P1(config-if)# ip address 10.12.96.2/32 secondary
```
2. Configure the global TLDP parameters including the router ID and transport address.
```
P1(config)#router ldp
P1(config-router)#router-id 10.12.96.2
P1(config-router)# targeted-peer ipv4 10.12.96.4
P1(config-router-targeted-peer)#exit
P1(config-router)# transport-address ipv4 10.12.96.2
```
3. Configure LDP to prefer tunneling over RSVP.
```
P1(config)#router ldp
P1(config-router)# prefer-tunnel-in-tunnel rsvp
```
4. Configure global RSVP parameters.
```
P1(config)#router rsvp
```
5. Configure a RSVP trunk towards the neighbor router (assuming 10.12.96.4 is the neighbor) and enable ldp-tunneling to allow tunneling LDP LSPs.
```
P1(config)# rsvp-trunk t1 ipv4
P1(config-trunk)#to 10.12.96.4
P1(config-trunk)#ldp-tunneling
```
6. Configure the interface facing the network side with an IP address, enable label switching, and enable LDP and RSVP.
   - For interface xe1:
```
P1(config)#interface xe1
P1(config-if)#ip address 10.1.1.2/24
P1(config-if)#label-switching
P1(config-if)#enable-ldp ipv4
```
   - For interface xe2:
```
P1(config)#interface xe2
P1(config-if)#ip address 20.1.1.1/24
P1(config-if)#label-switching
P1(config-if)#enable-rsvp
```
7. If using ISIS as the Interior Gateway Protoco (IGP), configure ISIS parameters including traffic engineering.
```
P1(config)#router isis ISIS-IGP
P1(config-router)#is-type level-1
P1(config-router)#metric-style wide
P1(config-router)#mpls traffic-eng router-id 10.12.96.2
P1(config-router)#mpls traffic-eng level-1
P1(config-router)#capability cspf
P1(config-router)#dynamic-hostname
P1(config-router)#net 49.0000.0000.0002.00
P1(config-router)#exit
```

Configure LDP Tunneling over RSVP on P2 Router

Follow the steps to configure the LDP tunneling over RSVP on P2 router:

1. Configure the loopback interface with an IP address.

```
P2(config)#interface lo
P2(config-if)# ip address 10.12.96.3/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
P2(config)#router ldp
P2(config-router)#router-id 10.12.96.3
P2(config-router)# transport-address ipv4 10.12.96.3
```

3. Configure global RSVP parameters.

```
P2(config)#router rsvp
```

4. Configure the interface facing the network side with an IP address, enable label switching, and enable RSVP.
   - For interface xe2:

```
P2(config)#interface xe2
P2(config-if)#ip address 20.1.1.2/24
P2(config-if)#label-switching
P2(config-if)#enable-rsvp
```

   - For interface xe3:

```
P2(config)#interface xe3
P2(config-if)#ip address 30.1.1.1/24
P2(config-if)#label-switching
P2(config-if)#enable-rsvp
```

5. If using ISIS as the Interior Gateway Protoco (IGP), configure ISIS parameters including traffic engineering.

```
P2(config)#router isis ISIS-IGP
P2(config-router)#is-type level-1
P2(config-router)#metric-style wide
P2(config-router)#mpls traffic-eng router-id 10.12.96.3
P2(config-router)#mpls traffic-eng level-1
P2(config-router)#capability cspf
P2(config-router)#dynamic-hostname
P2(config-router)#net 49.0000.0000.0003.00
P2(config-router)#exit
```

Configure LDP Tunneling over RSVP on P3 Router

Follow the steps to configure the LDP tunneling over RSVP on P3 router:

1. Configure the loopback interface with an IP address.

```
P3(config)#interface lo
P3(config-if)# ip address 10.12.96.4/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
P3(config)#router ldp
P3(config-router)#router-id 10.12.96.4
P3(config-router)# targeted-peer ipv4 10.12.96.2
P3(config-router-targeted-peer)#exit
P3(config-router)# transport-address ipv4 10.12.96.4
```

3. Configure global RSVP parameters.

```
P3(config)#router rsvp
```

4. Configure prefix lists.

```
P3(config)# ip prefix-list fec_list
P3(config-ip-prefix-list)# seq 5 permit 10.12.96.5/32
P3(config)# ip prefix-list peer_list
P3(config-ip-prefix-list)# seq 5 permit 10.12.96.2/32
```

5. Configure prefix lists to control label advertisement between peers.

```
P3(config)# router ldp
P3(config-router)# advertise-labels for fec_list to peer_list
```

6. Configure the interface facing the network side with an IP address, enable label switching, and enable RSVP.
   - For interface xe3:

```
P3(config)#interface xe3
P3(config-if)#ip address 30.1.1.2/24
P3(config-if)#label-switching
P3(config-if)#enable-rsvp
```

   - For interface xe4:

```
P3(config)#interface xe4
P3(config-if)#ip address 40.1.1.1/24
P3(config-if)#label-switching
P3(config-if)#enable-ldp ipv4
```

7. If using ISIS as the Interior Gateway Protoco (IGP), configure ISIS parameters including traffic engineering.

```
P3(config)#router isis ISIS-IGP
P3(config-router)#is-type level-1
P3(config-router)#metric-style wide
P3(config-router)#mpls traffic-eng router-id 10.12.96.4
P3(config-router)#mpls traffic-eng level-1
P3(config-router)#capability cspf
P3(config-router)#dynamic-hostname
P3(config-router)#net 49.0000.0000.0004.00
P3(config-router)#exit
```

Configure LDP Tunneling over RSVP on PE2 Router

Follow the steps to configure the LDP tunneling over RSVP on PE2 router:

1. configure the loopback interface with an IP address.

```
PE2(config)#interface lo
PE2(config-if)# ip address 10.12.96.5/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
PE2(config)#router ldp
PE2(config-router)#router-id 10.12.96.5
PE2(config-router)# transport-address ipv4 10.12.96.5
```

3. Configure the interface facing the network side with an IP address, enable label switching, and enable LDP.

```
PE2(config)#interface xe4
PE2(config-if)#ip address 40.1.1.2/24
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
```

4. If using ISIS as the Interior Gateway Protoco (IGP), configure ISIS parameters including traffic engineering.

```
PE2(config)#router isis ISIS-IGP
PE2(config-router)#is-type level-1
PE2(config-router)#metric-style wide
PE2(config-router)#mpls traffic-eng router-id 10.12.96.5
PE2(config-router)#mpls traffic-eng level-1
PE2(config-router)#capability cspf
PE2(config-router)#dynamic-hostname
PE2(config-router)#net 49.0000.0000.0005.00
PE2(config-router)#exit
```

Snippet Configuration on P1 Router

Follow the steps to configure the LDP tunneling over RSVP on P1 router using snippet:

```
P1#show running-config isis
!
!
router isis ISIS-IGP-100
 is-type level-1
 metric-style wide
 mpls traffic-eng router-id 10.12.96.2
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 net 49.0001.0000.0000.0002.00
!

P1#show running-config ldp
!
router ldp
 router-id 10.12.96.2
 prefer-tunnel-in-tunnel rsvp
 targeted-peer ipv4 10.12.96.4
  exit-targeted-peer-mode
 transport-address ipv4 10.12.96.2
!
```

```
interface xe1
 enable-ldp ipv4


P1#show running-config rsvp
!
router rsvp
!
!
interface xe2
 enable-rsvp
!
!
rsvp-trunk t1 ipv4
 to 10.12.96.4
 ldp-tunneling
!
```

Snippet Configuration on P3 Router

Follow the steps to configure the LDP tunneling over RSVP on P3 router using snippet:

```
P3#show running-config ldp
!
router ldp
 targeted-peer ipv4 10.12.96.2
  exit-targeted-peer-mode
 transport-address ipv4 10.12.96.4
 advertise-labels for fec_list to peer_list
!
interface xe4
 enable-ldp ipv4
!
```

Validation

```
Validation on P1 node:

P1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       10.12.96.1             xe1      Active    OPERATIONAL   30    00:05:42
       10.12.96.4              xe2      Passive   OPERATIONAL   30    00:05:44
P1#
P1#
P1#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to
Secondary
indicates the session is active with local repair at one or more nodes
```

```
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName
State Uptime    Rt  Style  Labelin  Labelout
10.12.96.4      10.12.96.2      5001     2201     PRI    t1-Primary                    UP
00:01:15  1 1 SE      -         25600
Total 1 displayed, Up 1, Down 0.

P1#
P1#
P1#
P1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled
Code    FEC               FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-Intf
ELC       Nexthop         UpTime
   L>   10.12.96.1/32      2        39        -          -     -           -          -
-               00:31:26
                                      38        -           Yes   3           xe1
No        10.1.1.1        -
R(t)>   10.12.96.4/32       1         9         5001        Yes   25600       xe2
No        22.1.1.1        00:01:19
   L>   10.12.96.5/32      3        11        -          -     -           -          -
-               00:01:19
                                   10        -          Yes   26244       -          No
10.12.96.4      -
                                                                                    (via rsvp tunnel-
id 5001, nhlfe_ix 9, label 25600)
P1#
P1#
P1#
P1#show ldp tunneling

Tunnel Name        : t1
Tunnel Endpoint    : 10.12.96.4/32
Tunnel Cost        : 20
Tunnel Owner       : RSVP
Tunnel Status      : Up

FEC                 Upstream-Peer   In-Label   Out-Label
=========================================================
10.12.96.5/32       10.12.96.1      26242     26244

Total FEC tunneld by t1 : 1

P1#
```

```
P1#
P1#show ldp tunneling-fec
FEC                Tunnel-name                    Tunnel-endpoint     Upstream-Peer   In-
label  Out-label
================================================================================================
==================
10.12.96.5/32      t1                        10.12.96.4/32      10.12.96.1     26242
26244

Total LDP Tunneled FEC : 1
P1#
P1#
P1#
P1#
P1#sh ldp tunneling-tunnels
Tunnel-name                   Tunnel-endpoint     Status    Cost
=========================================================================
t1                            10.12.96.4/32       Up        20
```

CLI Commands for LDP Tunneling over RSVP-TE

The LDP Tunneling over RSVP-TE introduces the following configuration commands.

ldp-tunneling

Use this command to enable LDP tunneling over RSVP trunk. When a specific RSVP trunk is enabled for tunneling, user traffic is tunneled using LDP LSP over RSVP LSP. If more than one trunk is enabled for tunneling LDP LSP, following trunk selection method is followed:

- If there are more than one trunk with same tunnel end-node, trunk with best metric (lower cost) is selected.
- If a destination FEC is reachable via more than one tunnel-endpoint, a tunnel-endpoint which is closer to destination is selected for tunneling.

Note: TLDP sessions should be manually established with RSVP tunnel end-nodes. Additionally, the 'advertise-labels' CLI must be explicitly configured to permit label advertisement over TLDP sessions.

Use `no` parameter of this command to disable tunneling from a trunk.

Command Syntax

```
ldp-tunneling
no ldp-tunneling
```

Parameters

None

Default

Disabled

Command Mode

rsvp-trunk mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to enable LDP tunneling over RSVP trunk:

```
OcNOS#configure terminal
OcNOS(config)#rsvp-trunk t2
OcNOS(config-trunk)#to 4.4.4.4
OcNOS(config-trunk)#ldp-tunneling
OcNOS(config-trunk)#commit
OcNOS(config-trunk)#end
```

prefer-tunnel-in-tunnel rsvp

Use this command for prioritizing RSVP trunk over LDP-LSP for forwarding LDP traffic. By default incoming LDP traffic is forwarded using LDP LSP. However when this CLI is configured and if RSVP trunk has been enabled for tunneling LDP LSP, user data (incoming LDP LSP) is tunneled over RSVP tunnels. If this CLI is not enabled and RSVP trunk has been enabled for tunneling LDP LSP, user data still can be forwarded over RSVP trunk if no LDP LSP exist.

Use `no` parameter of this command to prioritizing LDP-LSP over RSVP trunk while forwarding LDP traffic.

Command Syntax

```
prefer-tunnel-in-tunnel rsvp
```

```
no prefer-tunnel-in-tunnel rsvp
```

Parameters

None

Default

LDP-LSP is selected over RSVP trunks for forwarding.

Command Mode

Router LDP mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to prioritize RSVP trunk over LDP-LSP for forwarding LDP traffic:

```
OcNOS#configure terminal
OcNOS(config)#router ldp
OcNOS(config-router)#prefer-tunnel-in-tunnel rsvp
OcNOS(config-router)#commit
OcNOS(config-router)#end
```

Show Commands for LDP Over RVSP

show ldp tunneling fec

This command displays the LDP tunneling FEC mappings.

Command Syntax

```
show ldp tunneling-fec
```

Parameters

None

Command Mode

EXEC mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following configuration illustrates how to view the FEC mappings on router R2:

```
R2#show ldp tunneling-fec
FEC                 Tunnel-name                  Tunnel-endpoint
Upstream-Peer   In-label  Out-label
52.1.1.0/24         t2                          4.4.4.4/32          1.1.1.1
26253     26250
53.1.1.0/24         t2                          4.4.4.4/32          1.1.1.1
26255     26241

Total LDP Tunneled FEC : 2
```

show ldp tunneling

This command displays the LDP tunneling.

Command Syntax

```
show ldp tunneling
```

Parameters

None

Command Mode

EXEC mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to view the LDP tunneling on router R2:

```
Tunnel Name         : t1
```

```
Tunnel Endpoint    : 10.12.96.4/32
Tunnel Cost        : 20
Tunnel Owner       : RSVP
Tunnel Status      : Up

FEC                 Upstream-Peer   In-Label   Out-Label
============================================================
10.12.96.5/32       10.12.96.1      26242     26244

Total FEC tunneld by t1 : 1
```

show ldp tunneling-tunnels.

This command displays the LDP tunneling tunnels.

Command Syntax

```
show ldp tunneling-tunnels
```

Parameters

None

Command Mode

EXEC mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to view the LDP tunneling on router R2:

```
R2#show ldp tunneling-tunnels
Tunnel-name                   Tunnel-endpoint     Status    Cost
t2                            4.4.4.4/32          Up        20
```

Glossary

Note: List key terms used in this document and add the term and explanation to our existing Glossary.

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Forward Error Correction (FEC)	A system of error control that allows the receiver to correct some errors without having to request a re-transmission of data.
Interior Gateway Protocol (IGP)	An intradomain protocol used to exchange network reachability and routing information among devices within an autonomous system (AS), such as Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF), or Routing Information Protocol (RIP). Contrast with Exterior Gateway Protocol (EGP).

Label Distribution Protocol (LDP)	A protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to create label-switched path (LSP) instances through a network by mapping network layer routing information directly to data-link layer switched paths.
Resource Reservation Protocol (RSVP)	A signaling protocol for reserving resources across a network. RSVP is rarely used by itself, but Resource Reservation Protocol—Traffic Engineering (RSVP-TE) is widely used.
Targeted Label Distribution Protocol (TLDP)	A specialized form of LDP (Label Distribution Protocol) sessions.

CHAPTER 4 Hierarchical VPLS

Overview

A Virtual Private LAN Service (VPLS) enables multipoint to multipoint communication, creating LAN-like connectivity between customers' sites. However, the typical full mesh topology required for LAN emulation can be impractical in large networks. To address this, Hierarchical VPLS (H-VPLS) introduces a hierarchical approach using a spoke-PW (pseudowire) type. Unlike the standard mesh-PW, the spoke-PW facilitates traffic between hierarchical levels, offering a more scalable solution for VPLS networks.

H-VPLS Redundancy Characteristics

In a Virtual Private LAN Service (VPLS) network, when a node connects through a spoke-PW, a single point of failure arises. In the event of a connection failure to the VPLS mesh or a failure within the PE-rs node, the spoke device experiences a complete loss of connectivity. To address this, PW redundancy is implemented, configuring a secondary path that activates if the primary path fails. The MTU-s is configured with a primary spoke-PW connected to PE1-rs and a secondary spoke-PW connected to PE2-rs. During normal operation, the primary spoke-PW is active, but in case of failure, the MTU-s can switch to the standby spoke-PW for continued connectivity, aiming for sub-second convergence times with potential MAC flush-related traffic loss.

Benefits

Hierarchical VPLS (H-VPLS) is introduced to address scalability challenges associated with the traditional VPLS (Virtual Private LAN Service) architecture. It introduces a hierarchical approach that enhances scalability, reduces configuration complexity, optimizes traffic flow, and improves overall network efficiency and fault tolerance.

Limitations

- Automatic revertive cases from secondary to primary will not be supported.
- MAC Address Withdrawal feature will not be supported in release 6.5.2.
- Convergence on redundancy may require bidirectional traffic or MAC aging.

Prerequisites

- The `block-mesh-spoke-on-all-ac-down` and `ignore-ac-spoke-state` commands are optional and mutually exclusive, meaning only one can be applied at a time, or neither. By default, neither command is applied. If one of commands is applied, applying the other will make it the active one. To remove a command, use the `no` prefix.

```
signaling ldp
   (block-mesh-spoke-on-all-ac-down | ignore-ac-spoke-state)
   (no block-mesh-spoke-on-all-ac-down | no ignore-ac-spoke-state)
```

- **Define Interfaces and Loopback Addresses:**
  Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connec-tivity. These configurations establish the efficient network routing and communication.

```
!
interface lo
```

```
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128

interface xe14
 ip address 30.1.1.2/24
```

- **Configure IGP for Dynamic Routing:** Enable ISIS to facilitate dynamic routing on all nodes within the net-work. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution.Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

  ISIS Configuration:

```
router isis 1
 is-type level-2-only
 metric-style wide
 microloop-avoidance level-2
 mpls traffic-eng router-id 2.2.2.2
 mpls traffic-eng level-2
 capability cspf
 dynamic-hostname
 bfd all-interfaces
 net 49.0000.0000.0002.00
 passive-interface lo
!
interface xe14
 mpls ldp-igp sync isis level-2
 isis network point-to-point
 ip router isis 1
```

  OSPF Configuration:

```
router ospf 1
ospf router-id 2.2.2.2
network 2.2.2.2/32 area 0.0.0.0
 network 30.1.1.0/24 area 0.0.0.0!
!
interface xe14
ip ospf network point-to-point
```

Configuration for H-VPLS with Redundancy

Configure various nodes within the topology to set up a H-VPLS session.

Topology

This sample topology provides basic connectivity and routing between the devices.

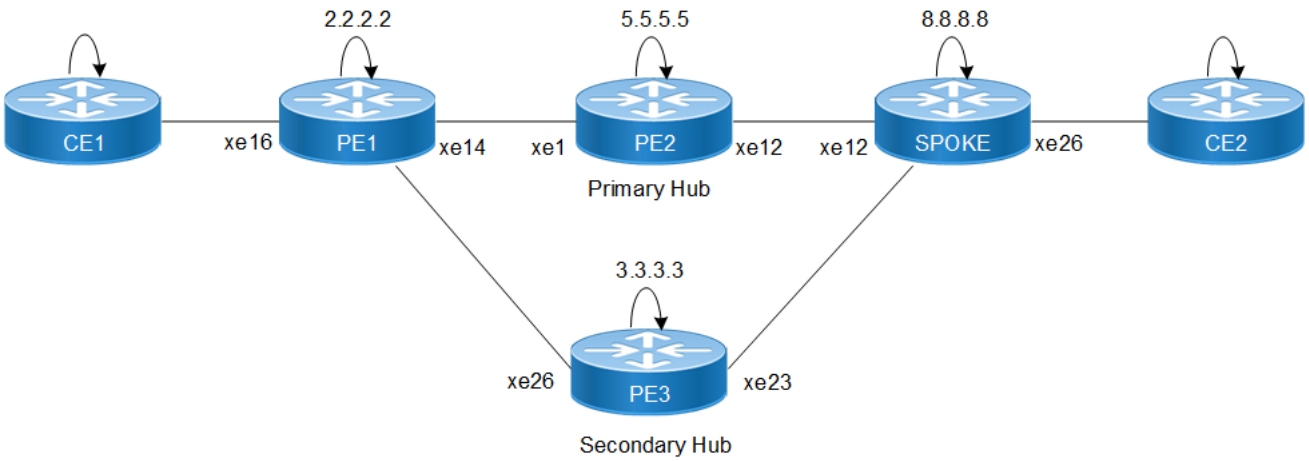

Figure 4-1: H-VPLS Configuration with Redundancy

Configure H-VPLS on PE1 Router

Follow the steps to configure the H-VPLS on PE1 router:

1. Configure router LDP.

```
PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2
```

2. Configure targeted-peer under router LDP.

```
PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)#  exit-targeted-peer-mode
PE1(config-router)# targeted-peer ipv4 3.3.3.3
PE1(config-router-targeted-peer)#  exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
PE1(config)#interface xe14
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

PE1(config)#interface xe26
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)#  vpls-peer 3.3.3.3
PE1(config-vpls-sig)#  vpls-peer 5.5.5.5
PE1(config-vpls-sig)#  exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
```

5. Configure sub-interface and attach vpls-instance to sub-interface.

```
PE1(config)#
PE1(config)#interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
```

```
PE1(config-acc-if-vpls)#  mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

Configure H-VPLS on PE2 (Primary Hub)

Follow the steps to configure the H-VPLS on PE2 (Primary Hub):

1. Configure router LDP.

```
PE2(config)#router ldp
PE2(config-router)# router-id 5.5.5.5
PE2(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
PE2(config)#router ldp
PE2(config-router)# targeted-peer ipv4 2.2.2.2
PE2(config-router-targeted-peer)#  exit-targeted-peer-mode
PE2(config-router)# targeted-peer ipv4 3.3.3.3
PE2(config-router-targeted-peer)#  exit-targeted-peer-mode
PE2(config-router)#
```

3. Enable LDP and label-switching for core interface.

```
PE2(config)#interface xe1
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching

PE2(config)#interface xe12
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE2(config)#mpls vpls vpls2000 2000
PE2(config-vpls)# signaling ldp
PE2(config-vpls-sig)#  vpls-peer 2.2.2.2
PE2(config-vpls-sig)#  vpls-peer 3.3.3.3
PE2(config-vpls-sig)#  exit-signaling
PE2(config-vpls)# exit-vpls
PE2(config)#
```

5. Configure L2-ckt.

```
PE2(config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
PE2(config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
PE2(config)#mpls vpls vpls2000 2000
PE2(config-vpls)#vpls-vc vc2000
PE2(config-vpls-spoke)#
```

Configure H-VPLS on PE3 (Secondary Hub)

Follow the steps to configure the H-VPLS on PE3 (Secondary Hub):

1. Configure router LDP.

```
PE3(config)#router ldp
PE3(config-router)# router-id 3.3.3.3
PE3(config-router)# transport-address ipv4 3.3.3.3
```

2. Configure targeted-peer under router LDP.

```
PE3(config)#router ldp
PE3(config-router)# targeted-peer ipv4 2.2.2.2
PE3(config-router-targeted-peer)#  exit-targeted-peer-mode
PE3(config-router)# targeted-peer ipv4 5.5.5.5
PE3(config-router-targeted-peer)#  exit-targeted-peer-mode
PE3(config-router)#
```

3. Enable LDP and label-switching for core interface.

```
PE3(config)#interface xe23
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching

PE3(config)#interface xe26
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE3(config)#mpls vpls vpls2000 2000
PE3(config-vpls)# signaling ldp
PE3(config-vpls-sig)#  vpls-peer 2.2.2.2
PE3(config-vpls-sig)#  vpls-peer 5.5.5.5
PE3(config-vpls-sig)#  exit-signaling
PE3(config-vpls)# exit-vpls
PE3(config)#
```

5. Configure L2-ckt.

```
PE3(config)#mpls l2-circuit vc2001 2223 8.8.8.8 mode raw
PE3(config-pseudowire)#
```

6. 6.Attach L2-ckt under vpls instance.

```
PE3 (config)#mpls vpls vpls2000 2000
PE3(config-vpls)#vpls-vc vc2001
PE3(config-vpls-spoke)#
```

Configure H-VPLS on Spoke Router

Follow the steps to configure the H-VPLS on Spoke router:

1. Configure router LDP.

```
Spoke(config)#router ldp
Spoke(config-router)# router-id 8.8.8.8
Spoke(config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke(config-router)# targeted-peer ipv4 5.5.5.5
Spoke(config-router-targeted-peer)#  exit-targeted-peer-mode
Spoke(config-router)# targeted-peer ipv4 3.3.3.3
Spoke(config-router-targeted-peer)#  exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke(config)#interface xe12
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching

Spoke(config)#interface xe25
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

4. Configure VPLS instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#
```

5. 5.Configure L2-ckt.

```
Spoke(config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke(config-pseudowire)#!
Spoke(config-pseudowire)#mpls l2-circuit vc2001 2223 3.3.3.3 mode raw
Spoke(config-pseudowire)#
```

6. 6.Configure Primary and secondary spoke under vpls instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#vpls-vc vc2000
Spoke(config-vpls-spoke)#  secondary vc2001
Spoke(config-vpls-spoke)#  exit-spoke
Spoke(config-vpls)# exit-vpls
Spoke(config)#
```

7. Configure sub-interface and attach vpls-instance to sub-interface.

```
Spoke(config)#
Spoke(config)#interface xe26.2000 switchport
Spoke(config-if)# encapsulation dot1q 2000
Spoke(config-if)# access-if-vpls
Spoke(config-acc-if-vpls)#  mpls-vpls vpls2000
Spoke(config-acc-if-vpls)#
```

Running Configuration on PE1 Router

```
router ldp
 router-id 2.2.2.2
 targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
 targeted-peer ipv4 5.5.5.5
 transport-address ipv4 2.2.2.2
!
interface xe14
     enable-ldp ipv4
!
interface xe26
     enable-ldp ipv4
!
mpls vpls vpls2000 2000
 signaling ldp
 vpls-peer 3.3.3.3
 vpls-peer 5.5.5.5
  exit-signaling
 exit-vpls
!
interface xe16.2000 switchport
 access-if-vpls
    mpls-vpls vpls2000
```

Running Configuration on PE2 Router

```
router ldp
 targeted-peer ipv4 2.2.2.2
```

```
  exit-targeted-peer-mode
 targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
transport-address ipv4 5.5.5.5
!
mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
 vpls-vc vc2000
  exit-spoke
 signaling ldp
  vpls-peer 2.2.2.2
  vpls-peer 3.3.3.3
  exit-signaling
 exit-vpls
```

Running Configuration on PE3 Router

```
router ldp
 targeted-peer ipv4 2.2.2.2
  exit-targeted-peer-mode
 targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
transport-address ipv4 3.3.3.3
!
mpls l2-circuit vc2001 2223 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
 vpls-vc vc2001
  exit-spoke
 signaling ldp
  vpls-peer 2.2.2.2
  vpls-peer 5.5.5.5
  exit-signaling
 exit-vpls
```

Running Configuration on Spoke Router

```
router ldp
 router-id 8.8.8.8
 targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
 targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
 transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls l2-circuit vc2001 2223 3.3.3.3 mode raw
!
mpls vpls vpls2000 2000
vpls-vc vc2000
  secondary vc2001
  exit-spoke
 exit-vpls
!
```

```
interface xe26.2000 switchport
 access-if-vpls
    mpls-vpls vpls2000
```

Validation

Validate the show output after configuration as shown below.
Verify vpls mesh are up between PE1 and Hub Nodes

```
PE1#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr         Tunnel-Label  In-Label   Network-Intf   Out-Label
Lkps/St   PW-INDEX SIG-Protocol   Status     UpTime
2000       3.3.3.3           29447         28164      xe26           27532
2/Up      3         LDP           Active     2d12h08m
2000       5.5.5.5           31364         28162      xe14           26883
2/Up      4         LDP           Active     2d12h04m

PE2#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr         Tunnel-Label  In-Label   Network-Intf   Out-Label
Lkps/St   PW-INDEX SIG-Protocol   Status     UpTime
2000       2.2.2.2           29446         26883      xe1        28162      2/
Up      3         LDP           Active     2d12h05m
2000       3.3.3.3           31367         26884      xe1            27528
2/Up      4         LDP           Active     2d12h15m

PE3#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr         Tunnel-Label  In-Label   Network-Intf   Out-Label
Lkps/St   PW-INDEX SIG-Protocol   Status     UpTime
2000       2.2.2.2           29440         27532      xe26           28164
2/Up      3         LDP           Active     2d12h10m
2000       5.5.5.5           31363         27528      xe26           26884
2/Up      4         LDP           Active     2d12h16m


Verify vpls spoke between Hub and Spoke

PE2#sho mpls vpls spoke
VPLS-ID    Virtual Circuit  Tunnel-Label In-Label   Network-Intf   Out-Label
Lkps/St   Secondary
2000       vc2000           29443         26882      xe1            26886
2/Up     ---

PE3#show mpls vpls spoke
VPLS-ID    Virtual Circuit  Tunnel-Label In-Label   Network-Intf   Out-Label
Lkps/St   Secondary
2000       vc2001           N/A           27527      N/A            26883
0/Dn     ---

Spoke#show mpls vpls spoke
```

```
VPLS-ID    Virtual Circuit  Tunnel-Label In-Label   Network-Intf   Out-Label
Lkps/St   Secondary
2000       vc2000           29440         26886      xe12           26882
2/Up     vc2001
2000       vc2001           N/A           26883      N/A            27527
0/Dn     ---
```

Verify H-vpls session on Hub and spoke:

```
PE2#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
 SIG-Protocol: LDP
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
 Description: none
 service-tpid: dot1.q
 Operating mode: Raw
 Ignoring AC interface and spoke-VC state

 Configured interfaces:
   None

 Mesh Peers:
   2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h13m)
   3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h22m)
 Spoke Peers:
   vc2000 (Up) (UpTime 01:31:27)


PE3#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
 SIG-Protocol: LDP
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
 Description: none
 service-tpid: dot1.q
 Operating mode: Raw
 Ignoring AC interface and spoke-VC state

 Configured interfaces:
   None

 Mesh Peers:
   2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h16m)
   5.5.5.5 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h22m)
 Spoke Peers:
   vc2001 (Dn) (Reason: VC on standby)


Spoke#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
```

```
SIG-Protocol: N/A
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw

Configured interfaces:
 Interface: xe26.2000
  Status: Up
  Subinterface Match Criteria(s) :
  dot1q 2000

Spoke Peers:
  vc2000 (Up) (UpTime 01:31:33)
   Secondary: vc2001 (Dn) (Reason: VC on standby)
```

Configuration for H-VPLS without Redundancy

Configure various nodes within the topology to set up a H-VPLS session.

Topology

This sample topology provides basic connectivity and routing between the devices.

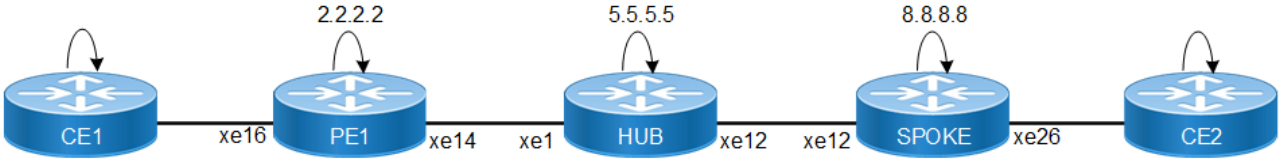

Figure 4-2: H-VPLS Configuration without Redundancy

Configure H-VPLS on PE1 Router

Follow the steps to configure the H-VPLS on PE1 router:

1. Configure router LDP.

```
PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2
```

2. Configure targeted-peer under router LDP.

```
PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)#  exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
PE1(config)#interface xe14
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)#  vpls-peer 5.5.5.5
PE1(config-vpls-sig)#  exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
```

5. Configure sub-interface and attach vpls-instance to sub-interface

```
PE1(config)#
PE1(config)#interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)#  mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

Configure H-VPLS on Hub Router

Follow the steps to configure the H-VPLS on Hub router:

1. Configure router LDP.

```
Hub(config)#router ldp
Hub(config-router)# router-id 5.5.5.5
Hub(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
Hub(config-router)# targeted-peer ipv4 2.2.2.2
Hub(config-router-targeted-peer)#  exit-targeted-peer-mode
R5-P5(config-router)# targeted-peer ipv4 8.8.8.8
R5-P5(config-router-targeted-peer)#
```

3. Enable LDP and label-switching for core interface.

```
Hub(config)#interface xe1
Hub(config-if)# enable-ldp ipv4
Hub(config-if)#label-switching

Hub(config)#interface xe12
Hub(config-if)# enable-ldp ipv4
Hub(config-if)#label-switching
```

4. Configure VPLS instance.

```
Hub(config)#mpls vpls vpls2000 2000
Hub(config-vpls)# signaling ldp
Hub(config-vpls-sig)# vpls-peer 2.2.2.2
Hub(config-vpls-sig)# exit-signaling
Hub(config-vpls)# exit-vpls
Hub(config)#
```

5. Configure L2-ckt.

```
Hub (config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
Hub (config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
Hub (config)#mpls vpls vpls2000 2000
Hub (config-vpls)#vpls-vc vc2000
Hub(config-vpls-spoke)#
```

Configure H-VPLS on Spoke Router

Follow the steps to configure the H-VPLS on Spoke router:

1. Configure router LDP.

```
Spoke(config)#router ldp
Spoke(config-router)# router-id 8.8.8.8
Spoke(config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke(config-router)# targeted-peer ipv4 5.5.5.5
Spoke(config-router-targeted-peer)#  exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke(config)#interface xe12
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

4. Configure VPLS instance.

```
Spoke(config)#mpls vpls vpls2000 2000
     Spoke(config-vpls)#
```

5. Configure L2-ckt.

```
Spoke(config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke(config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#vpls-vc vc2000
Spoke(config-vpls-spoke)#
```

7. Configure sub-interface and attach vpls-instance to sub-interface.

```
Spoke(config)#
Spoke(config)#interface xe26.2000 switchport
Spoke(config-if)# encapsulation dot1q 2000
Spoke(config-if)# access-if-vpls
Spoke(config-acc-if-vpls)#  mpls-vpls vpls2000
Spoke(config-acc-if-vpls)#
```

Running Configuration on PE1 Router

```
router ldp
 router-id 2.2.2.2
 targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
 transport-address ipv4 2.2.2.2
!
interface xe14
     enable-ldp ipv4
!
mpls vpls vpls2000 2000
 signaling ldp
  vpls-peer 5.5.5.5
  exit-signaling
 exit-vpls
!
interface xe16.2000 switchport
```

```
 access-if-vpls
    mpls-vpls vpls2000
```

Running Configuration on Hub Router

```
router ldp
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
 targeted-peer ipv4 8.8.8.8
  exit-targeted-peer-mode
!
!
mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
 vpls-vc vc2000
  exit-spoke
 signaling ldp
  vpls-peer 2.2.2.2
  exit-signaling
 exit-vpls
```

Running Configuration on Spoke Router

```
router ldp
 router-id 8.8.8.8
 targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
 transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls vpls vpls2000 2000
 vpls-vc vc2000
  exit-spoke
 exit-vpls
!
interface xe26.2000 switchport
 access-if-vpls
    mpls-vpls vpls2000
```

Validation

```
Validate the show output after configuration as shown below.
Verify vpls mesh are up between PE and Hub

PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr         Tunnel-Label  In-Label   Network-Intf   Out-Label
Lkps/St   PW-INDEX SIG-Protocol   Status     UpTime
2000       5.5.5.5           31364         28162      xe14           26883
2/Up      4          LDP           Active     2d10h36m
```

```
Hub#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr         Tunnel-Label  In-Label   Network-Intf   Out-Label
Lkps/St   PW-INDEX SIG-Protocol   Status     UpTime
2000       2.2.2.2           29446         26883          xe1        28162
2/Up      3         LDP           Active     2d10h39m


Verify vpls spoke are up between Hub and Spoke

Hub#sho ldp mpls-l2-circuit
Transport   Client     VC       VC           Local     Remote    Destination
Lo-cal                          Remote
VC ID       Binding    State    Type         VC Label  VC Label  Address
PW Status                       PW Status
2222        VPLS:2000  UP       Ethernet      26882     26886     8.8.8.8
Forwarding                      Forwarding

Hub#sho mpls vpls spoke
VPLS-ID    Virtual Circuit  Tunnel-Label In-Label   Network-Intf   Out-Label
Lkps/St   Secondary
2000       vc2000           29443         26882      ce4            26886
2/Up
    ---

Spoke#show ldp mpls-l2-circuit
Transport   Client     VC       VC           Local     Remote    Destination
Lo-cal                          Remote
VC ID       Binding    State    Type         VC Label  VC Label  Address
PW Status                       PW Status
2222        VPLS:2000  UP       Ethernet      26886     26882     5.5.5.5
Forwarding                      Forwarding

Spoke#show mpls vpls spoke
VPLS-ID    Virtual Circuit  Tunnel-Label In-Label   Network-Intf   Out-Label
Lkps/St   Secondary
2000       vc2000           29440         26886      ce4            26882
2/Up     ---

Verify H-vpls session on Hub and spoke:

Hub#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
 SIG-Protocol: LDP
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
 Description: none
 service-tpid: dot1.q
 Operating mode: Raw
 Ignoring AC interface and spoke-VC state

 Configured interfaces:
   None
```

```
 Mesh Peers:
   2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d10h47m)
   3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d10h56m)
 Spoke Peers:
   vc2000 (Up) (UpTime 00:05:48)


Spoke#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
 SIG-Protocol: N/A
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, Configured MTU: 1500
 Description: none
 service-tpid: dot1.q
 Operating mode: Raw

 Configured interfaces:
  Interface: xe26.2000
   Status: Up
   Subinterface Match Criteria(s) :
   dot1q 2000

 Spoke Peers:
   vc2000 (Up) (UpTime 00:07:47)
```

Commands for H-VPLS Configuration

The H-VPLS uses the following configuration commands.

vpls-vc

Use this command to add a spoke virtual circuit to VPLS domain hierarchically.

Use `no` parameter of this command to remove this configuration.

Command Syntax

```
vpls-vc NAME
 (secondary NAME|)
 (ethernet|vlan|)
```

Parameters

`NAME`	Specifies the name of the VPLS. It is a string that identifies the MPLS VC to add to the VPLS domain.
`secondary`	Specifies the name of the secondary spoke.
`NAME`	Specifies the name for the secondary spoke.
`ethernet`	Specifies the spoke type. Defaults to `ethernet`.
`vlan`	Specifies the spoke type.

Default

Disabled

Command Mode

VPLS mode

Applicability

Introduced before OcNOS version 1.3.

Modified the command prompt into a hierarchical structure from single line in the OcNOS version 6.5.1.

Example

Example for adding a spoke virtual circuit with VPLS name vc1 and secondary spoke vc2:

```
#configure terminal
(config)#mpls vpls vpls1 3000
(config-vpls)#vpls-vc vc1
(config-vpls-spoke)#secondary vc2
(config-vpls-spoke)#type ethernet
(config-vpls-spoke)#exit-spoke
(config-vpls)#exit
```

Example to remove the configuration of the spoke virtual circuit with VPLS name vc1:

```
#configure terminal
(config)#mpls vpls vpls1 3000
(config-vpls)#no vpls-vc vc1
(config-vpls)#exit
```

signaling

Use this command to set all mesh and spoke pseudowires to down when all access interfaces are down.

Use `ignore-ac-spoke-state` parameter of this command to remove this configuration.

Command Syntax

```
signaling ldp block-mesh-spoke-on-all-ac-down
signaling ignore-ac-spoke-state
```

Parameters

`block-mesh-spoke-on-all-ac-down`	(Optional) Controls the behavior of pseudowires (PWs) in a VPLS instance when all access interfaces associated with the VPLS instance are down.
`ignore-ac-spoke-state`	Ignores access interfaces and spoke pseudowires state and keep mesh pseudowires up.

Default

disabled

Command Mode

VPLS mode

Applicability

Introduced before OcNOS version 1.3.

Modified the command prompt into a hierarchical structure from single line in the OcNOS version 6.5.1.

Example

Example for setting up all mesh and spoke pseudowires to down when all access interfaces are down:

```
#configure terminal
(config)# mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#block-mesh-spoke-on-all-ac-down
(config-vpls-sig)#exit
```

Example for setting up all mesh and spoke pseudowires to up:

```
#configure terminal
(config)# mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#ignore-ac-spoke-state
(config-vpls-sig)#exit
```

CHAPTER 5 Auto-Bandwidth with RSVP-TE

Overview

Automatic bandwidth allows to dynamically adjust bandwidth reservation based on the measured traffic. RSVP automatic bandwidth monitors the traffic rate on a Label Switched Path (LSP) and resizes the bandwidth to align it closely with the traffic in the tunnel. RSVP automatic bandwidth is configured on individual LSPs at every headend router.

Auto bandwidth can be added to an operational LSP at any time, but no bandwidth change occurs until a future trigger event or auto bandwidth profile configured with initial bandwidth or minimum bandwidth. Auto bandwidth may also be removed from an operational LSP at any time and this would re-signal the LSP with no bandwidth reservation.

Feature Characteristics

The characteristics of the RSVP auto-bandwidth are:

- RSVP-TE auto bandwidth provides the means to automatically adjust the bandwidth allocation for traffic engineering tunnels based on their measured traffic load.
- This feature samples the average output rate for each tunnel marked for automatic bandwidth adjustment. For each marked tunnel, this feature periodically adjusts the tunnel's allocated bandwidth to the largest eligible sample for the tunnel since the last adjustment.
- The frequency with which tunnel bandwidth is adjusted and the allowable range of adjustments should be configurable on a per-auto-bandwidth profile basis.
- In addition, the sampling interval and the interval over which to average tunnel traffic to obtain the average output rate is user-configurable on a per-auto-bandwidth profile basis.

Note:

- Convergence on redundancy may require bidirectional traffic or MAC aging.
- The feature relies on stat_id allocation to tunnel entities, and there is a limit on the maximum number of stat_ids (which varies based on the chip variant). If a tunnel is not associated with a stat entity, traffic rate samples cannot be fetched for those tunnels.
- RSVP Graceful Restart is not supported for automatic bandwidth. When a GR is performed, RSVP will not store the current bandwidth for the reservation. It will use either the initial bandwidth (if configured), the minimum bandwidth, or the highest bandwidth of the on-boot sample (if auto-bandwidth-on-boot is configured).
- The auto bandwidth feature relies on the hardware's ability to collect tunnel traffic counters. In Qumran1 devices, the “hardware-profile statistics tunnel-lif enable” command must be enabled, and the system must be reloaded for the change to take effect. Without the tunnel statistics profile, auto bandwidth will not process traffic rates and will be ineffective. Note that, only 2 statistics profiles shall be configured as this is the hardware limitation.
- Auto bandwidth and manual bandwidth configurations are mutually exclusive. Auto bandwidth allows for configuring an initial bandwidth, which will be used as the session's initial bandwidth when auto bandwidth is associated with a trunk. If the initial bandwidth is not configured, the minimum bandwidth will be used to initialize the session bandwidth.
- For tunnels with only one hop, the no PHP (default config) must be set for the rate to be computed correctly.

Benefits

In large MPLS transport networks in service provider settings with this capability:

- The network can react faster to sudden bursts of traffic in near real-time and not rely on manual intervention.
- Effective use of bandwidth resources by minimizing the over-subscription/padding of LSP bandwidth.
- Maximizes the usage of available bandwidth and optimizes the network effectively to use preferred, shorter latency, paths first.

Prerequisites

Define Interfaces and Loopback Addresses

Configure Layer 3 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!
interface lo
 ip address 127.0.0.1/8
 ip address 100.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface xe1
 ip address 1.1.1.1/24
!
```

Configure IGP for Dynamic Routing

Configure IGP for dynamic routing by following the steps mentioned. This setup includes enabling ISIS for dynamic routing and configuring OSPF for the network.

ISIS Configuration

1. **Enable ISIS on all nodes**: Ensure that ISIS is running across the network to facilitate dynamic routing.
2. **Define ISIS Router Instances**: Set up instances to match loopback IP addresses.
3. **Add Network Segments to ISIS Areas**: This ensures proper route distribution.
4. **Set up Neighbor Relationships**: Use loopback IP addresses to establish these relationships for efficient route advertisement and convergence.

```
!
router isis 1
 is-type level-2-only
 metric-style wide
 mpls traffic-eng router-id 100.1.1.1
 mpls traffic-eng level-2
 capability cspf
 dynamic-hostname
 fast-reroute ti-lfa level-2 proto ipv4
 net 49.0000.0000.0001.00
 passive-interface lo
!
interface xe1
 isis network point-to-point
```

```
 ip router isis 1
!
```

OSPF Configuration

1. **Configure OSPF Router ID**: Assign a unique router ID for OSPF operations.
2. **Define OSPF Networks**: Include the loopback IP and other network segments in the OSPF area for route distribution.

```
!
router ospf 100
 ospf router-id 100.1.1.1
 network 100.1.1.1/32 area 0.0.0.0
 network 1.1.1.1/24 area 0.0.0.0
!
```

Configure RSVP for Efficient Network Operation

Enable Resource Reservation Protocol (RSVP) on all nodes to optimize traffic routing and quality of service. RSVP reserves network resources along specified paths to enhance network performance and reliability.

```
!
router rsvp
!
interface xe1
 label-switching
 enable-rsvp
!
```

Configure the RSVP Primary Path and Trunk

Establish a trunk is required on edge routers participating in label-switching using defined path. Configuring the RSVP path is optional.

```
!
rsvp-path PE1-PE2-1 mpls
 1.1.1.2 strict
 1.1.2.1 strict
!
rsvp-trunk PE1-PE2 ipv4
 reoptimize
 primary fast-reroute protection facility
 primary fast-reroute node-protection
 primary path PE1-PE2-1
 from 100.1.1.1
 to 100.1.1.3
!
```

Configuration for RSVP Auto-Bandwidth

Configure various nodes within the topology to set up a RSVP-Auto bypass tunnels.

Topology

The sample topology includes Edge Nodes (PE1 and PE2) and core Nodes (P1).

Primary path is defined via PE1-P1-PE2.

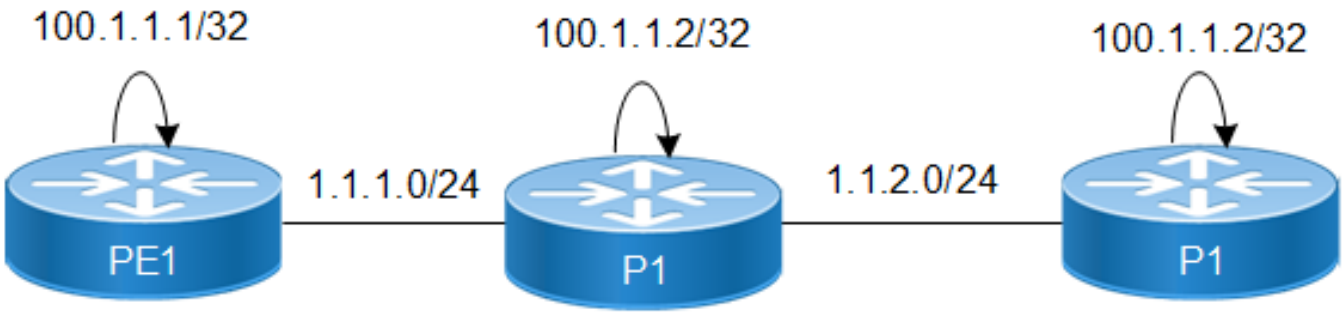

Figure 5-3: RSVP-Auto Bypass Tunnel Setup

Configure RSVP Auto Bandwidth on PE1 Router

1. Create auto-bandwidth Profile.

```
(config)# rsvp-auto-bandwidth AUTO-BW
(config-auto-bandwidth)# commit
```

2. Set the Sample interval & adjust interval.

```
(config-auto-bandwidth)# sample-interval 1
(config-auto-bandwidth)# adjust-interval 5
```

3. Set the minimum & maximum bandwidth rate.

```
(config-auto-bandwidth)# minimum-bandwidth 200m
(config-auto-bandwidth)# maximum-bandwidth 500m
```

4. Set the overflow-threshold & underflow-threshold.

```
(config-auto-bandwidth)# overflow-threshold absolute 100m
(config-auto-bandwidth)# underflow-threshold absolute 50m
```

5. Set the overflow & underflow limit.

```
(config-auto-bandwidth)# overflow-limit 2
(config-auto-bandwidth)# underflow-limit 2
```

6. Set the maximum number of consecutive times the average bandwidth can exceed the maximum threshold bandwidth before the exceed action is applied.

```
(config-auto-bandwidth)# maximum-bandwidth-exceed-limit 2
(config-auto-bandwidth)#maximum-bandwidth-exceed-action teardown
(config-auto-bandwidth)#commit
```

Running configuration on PE1 router is as follows:

```
#show running-config rsvp
!
router rsvp
!
!
interface xe1
 enable-rsvp
!
interface xe2
 enable-rsvp
!
```

```
!
!
!
rsvp-auto-bandwidth AUTO-BW
 sample-interval 1
 adjust-interval 5
 minimum-bandwidth 200m
 maximum-bandwidth 500m
 overflow-threshold absolute 100m
 underflow-threshold absolute 50m
 overflow-limit 2
 underflow-limit 2
 maximum-bandwidth-exceed-limit 2
 maximum-bandwidth-exceed-action teardown
!
rsvp-trunk PE1-PE2 ipv4
 reoptimize
 primary fast-reroute protection facility
 primary fast-reroute node-protection
 primary path PE1-PE2-1
 auto-bandwidth AUTO-BW
 from 100.1.1.1
 to 100.1.1.3
!
```

Validation

Verify auto bandwidth adjustments information as below:

```
Send the sample rate with 200MBPS and verify the auto bandwidth adjustments as below:

#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2202, Egress: 100.1.1.3
-------------------------------------------------------------------------------
Sample Interval                          : 1 minutes
Adjustment Interval                      : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                 : 0
Minimum Bandwidth                        : 200m
Maximum Bandwidth                        : 500m
Overflow Threshold Bandwidth             : 100m
Underflow Threshold Bandwidth            : 50m
Overflow Threshold Activate Bandwidth    : 0
Underflow Threshold Activate Bandwidth   : 0
Overflow Limit                           : 2
Underflow Limit                          : 2
Max. Bandwidth Exceed Limit              : 2
-------------------------------------------------------------------------------
Max-BW-exceed-limit action               : teardown
Resignal-failure-action                  : notify
Monitor Bandwidth                        : No
-------------------------------------------------------------------------------
```

```
Minimum Average Bandwidth               : 0
Maximum Average Bandwidth               : 202.8m
Total Overflow Count                    : 0
Consecutive Overflow Count              : 0
Consecutive Eligible Overflow Count     : 0
Total Underflow Count                   : 0
Consecutive Underflow Count             : 0
Consecutive Eligible Underflow Count    : 0
Max. Bandwidth Exceed Count             : 0
Teardown Count                          : 0
-------------------------------------------------------------------------------
Last Bandwidth                          : 0
Last Requested Bandwidth                : 0
Last Signaled Bandwidth                 : 0
Current Bandwidth                       : 200m
Highest Bandwidth                       : 203m
-------------------------------------------------------------------------------
Time for Next Sample request            : 48 seconds
Time for Next Adjustment                : 2 minutes, 57 seconds
Time of Last Bandwidth Request          : N/A
Time of Last Bandwidth Signal           : N/A
Time of Last Adjustment                 : N/A
Time of Highest Bandwidth Marked        : 2024 Jun 25 09:56:19
-------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments        : 0
Successful Adjustments                  : 0
Failed Adjustments                      : 0
-------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
  [Sample 1-5]       : 202.8m   202.6m


#show rsvp trunk auto-bandwidth
--------------------+--------+--------+--------+----------+----------+----------+------
----+------------+--------------------
Trunk-Name          Trunk    LSP      Last     Requested  Signaled   Current    Highest
Adjust-Time  Last-Adjust
                    ID       ID       BW       BW         BW         BW         BW
Left(sec)    Time
--------------------+--------+--------+--------+----------+----------+----------+------
----+------------+--------------------
PE1-PE2     5002     2202     0       0         0         200m      203m      176
NA


Overflow:

Current bandwidth is adjusted to 290.4MBPS. Then, send the sample rate which is more
than overflow threshold i.e, 340.1MBPS , 377.2MBPS.
As per the below output current bandwidth is more than overflow threshold bandwidth and
consecutively two samples are received and it is more than the overflow limit.
```

```
So current BW is adjusted to 377.2mbps after 2 consecutive samples collected as per the
Maximum Average Bandwidth.

#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2202, Egress: 100.1.1.3
-------------------------------------------------------------------------------
Sample Interval                          : 1 minutes
Adjustment Interval                      : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                 : 200m
Minimum Bandwidth                        : 100m
Maximum Bandwidth                        : 500m
Overflow Threshold Bandwidth             : 40m
Underflow Threshold Bandwidth            : 30m
Overflow Threshold Activate Bandwidth    : 0
Underflow Threshold Activate Bandwidth   : 0
Overflow Limit                           : 2
Underflow Limit                          : 2
Max. Bandwidth Exceed Limit              : 2
-------------------------------------------------------------------------------
Max-BW-exceed-limit action               : teardown
Resignal-failure-action                  : notify
Monitor Bandwidth                        : No
-------------------------------------------------------------------------------
Minimum Average Bandwidth                : 0
Maximum Average Bandwidth                : 340.1m
Total Overflow Count                     : 1
Consecutive Overflow Count               : 1
Consecutive Eligible Overflow Count      : 1
Total Underflow Count                    : 0
Consecutive Underflow Count              : 0
Consecutive Eligible Underflow Count     : 0
Max. Bandwidth Exceed Count              : 0
Teardown Count                           : 0
-------------------------------------------------------------------------------
Last Bandwidth                           : 190.7m
Last Requested Bandwidth                 : 290.4m
Last Signaled Bandwidth                  : 290.4m
Current Bandwidth                        : 290.4m
Highest Bandwidth                        : 340.1m
-------------------------------------------------------------------------------
Time for Next Sample request             : 0 seconds
Time for Next Adjustment                 : 0 seconds
Time of Last Bandwidth Request           : 2024 Jun 25 10:28:32
Time of Last Bandwidth Signal            : 2024 Jun 25 10:28:32
Time of Last Adjustment                  : 2024 Jun 25 10:28:32
Time of Highest Bandwidth Marked         : 2024 Jun 25 10:29:35
-------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments         : 4
Successful Adjustments                   : 4
Failed Adjustments                       : 0
```

```
-------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
  [Sample 1-5]       : 340.1m



#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2_1-Primary, Tunnel-id: 5002, LSP-ID: 2203, Egress: 100.1.1.3
-------------------------------------------------------------------------------
Sample Interval                         : 1 minutes
Adjustment Interval                     : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                : 200m
Minimum Bandwidth                       : 100m
Maximum Bandwidth                       : 500m
Overflow Threshold Bandwidth            : 40m
Underflow Threshold Bandwidth           : 30m
Overflow Threshold Activate Bandwidth   : 0
Underflow Threshold Activate Bandwidth  : 0
Overflow Limit                          : 2
Underflow Limit                         : 2
Max. Bandwidth Exceed Limit             : 2
-------------------------------------------------------------------------------
Max-BW-exceed-limit action              : teardown
Resignal-failure-action                 : notify
Monitor Bandwidth                       : No
-------------------------------------------------------------------------------
Minimum Average Bandwidth               : 0
Maximum Average Bandwidth               : 0
Total Overflow Count                    : 0
Consecutive Overflow Count              : 0
Consecutive Eligible Overflow Count     : 0
Total Underflow Count                   : 0
Consecutive Underflow Count             : 0
Consecutive Eligible Underflow Count    : 0
Max. Bandwidth Exceed Count             : 0
Teardown Count                          : 0
-------------------------------------------------------------------------------
Last Bandwidth                          : 290.4m
Last Requested Bandwidth                : 377.2m
Last Signaled Bandwidth                 : 377.2m
Current Bandwidth                       : 377.2m
Highest Bandwidth                       : 377.2m
-------------------------------------------------------------------------------
Time for Next Sample request            : 59 seconds
Time for Next Adjustment                : 0 seconds
Time of Last Bandwidth Request          : 2024 Jun 25 10:30:42
Time of Last Bandwidth Signal           : 2024 Jun 25 10:30:42
Time of Last Adjustment                 : 2024 Jun 25 10:30:42
Time of Highest Bandwidth Marked        : 2024 Jun 25 10:30:42
```

```
-------------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments        : 5
Successful Adjustments                  : 5
Failed Adjustments                      : 0
-------------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
=====================================
Underflow:

Scenario 1 :
Current bandwidth is adjusted to 377.2 mbps. Then, send the sample rate which is less
than underflow threshold i.e, 317.3 mbps , 310 mbps.
As per the below output current bandwidth is more than underflow threshold bandwidth and
consecutively two samples are received and it is more than the underflow limit.
So current BW is adjusted to 317.3 mbps after 2 consecutive samples collected as per the
Minimum Average Bandwidth.


#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2203, Egress: 100.1.1.3
-------------------------------------------------------------------------------------
Sample Interval                         : 1 minutes
Adjustment Interval                     : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                : 200m
Minimum Bandwidth                       : 100m
Maximum Bandwidth                       : 500m
Overflow Threshold Bandwidth            : 40m
Underflow Threshold Bandwidth           : 30m
Overflow Threshold Activate Bandwidth   : 0
Underflow Threshold Activate Bandwidth  : 0
Overflow Limit                          : 2
Underflow Limit                         : 2
Max. Bandwidth Exceed Limit             : 2
-------------------------------------------------------------------------------------
Max-BW-exceed-limit action              : teardown
Resignal-failure-action                 : notify
Monitor Bandwidth                       : No
-------------------------------------------------------------------------------------
Minimum Average Bandwidth               : 317.3m
Maximum Average Bandwidth               : 0
Total Overflow Count                    : 0
Consecutive Overflow Count              : 0
Consecutive Eligible Overflow Count     : 0
Total Underflow Count                   : 1
Consecutive Underflow Count             : 1
Consecutive Eligible Underflow Count    : 1
Max. Bandwidth Exceed Count             : 0
Teardown Count                          : 0
-------------------------------------------------------------------------------------
Last Bandwidth                          : 290.4m
```

```
Last Requested Bandwidth                : 377.2m
Last Signaled Bandwidth                 : 377.2m
Current Bandwidth                       : 377.2m
Highest Bandwidth                       : 377.2m
-------------------------------------------------------------------------------
Time for Next Sample request            : 9 seconds
Time for Next Adjustment                : 3 minutes, 6 seconds
Time of Last Bandwidth Request          : 2024 Jun 25 10:30:42
Time of Last Bandwidth Signal           : 2024 Jun 25 10:30:42
Time of Last Adjustment                 : 2024 Jun 25 10:30:42
Time of Highest Bandwidth Marked        : 2024 Jun 25 10:30:42
-------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments        : 5
Successful Adjustments                  : 5
Failed Adjustments                      : 0
-------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
  [Sample 1-5]       : 317.3m



#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2204, Egress: 100.1.1.3
-------------------------------------------------------------------------------
Sample Interval                         : 1 minutes
Adjustment Interval                     : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                : 200m
Minimum Bandwidth                       : 100m
Maximum Bandwidth                       : 500m
Overflow Threshold Bandwidth            : 40m
Underflow Threshold Bandwidth           : 30m
Overflow Threshold Activate Bandwidth   : 0
Underflow Threshold Activate Bandwidth  : 0
Overflow Limit                          : 2
Underflow Limit                         : 2
Max. Bandwidth Exceed Limit             : 2
-------------------------------------------------------------------------------
Max-BW-exceed-limit action              : teardown
Resignal-failure-action                 : notify
Monitor Bandwidth                       : No
-------------------------------------------------------------------------------
Minimum Average Bandwidth               : 0
Maximum Average Bandwidth               : 0
Total Overflow Count                    : 0
Consecutive Overflow Count              : 0
Consecutive Eligible Overflow Count     : 0
Total Underflow Count                   : 0
Consecutive Underflow Count             : 0
Consecutive Eligible Underflow Count    : 0
```

```
Max. Bandwidth Exceed Count             : 0
Teardown Count                          : 0
-------------------------------------------------------------------------------
Last Bandwidth                          : 377.2m
Last Requested Bandwidth                : 317.3m
Last Signaled Bandwidth                 : 317.3m
Current Bandwidth                       : 317.3m
Highest Bandwidth                       : 377.2m
-------------------------------------------------------------------------------
Time for Next Sample request            : 56 seconds
Time for Next Adjustment                : 2 minutes, 46 seconds
Time of Last Bandwidth Request          : 2024 Jun 25 10:32:55
Time of Last Bandwidth Signal           : 2024 Jun 25 10:32:55
Time of Last Adjustment                 : 2024 Jun 25 10:32:55
Time of Highest Bandwidth Marked        : 2024 Jun 25 10:30:42
-------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments        : 6
Successful Adjustments                  : 6
Failed Adjustments                      : 0
-------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
Scenario 2 :
Configure the Auto bandwidth Profile without configuring any underflow-limit. When all
the samples in adjustment cycle receive with the underflow rate, then only underflow
bandwidth adjustment will happen.
Below Example shows underflow limit as a Zero and current bandwidth is set to 8.5g and
all the samples are received less then underflow-limit. So, the Bandwidth adjustment
happens only after adjustment cycle.
#show rsvp trunk auto-bandwidth PE1-PE2
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2202, Egress: 100.1.1.3
-------------------------------------------------------------------------------
Sample Interval                         : 1 minutes
Adjustment Interval                     : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                : 4g
Minimum Bandwidth                       : 1g
Maximum Bandwidth                       : 9g
Overflow Threshold Bandwidth            : 10% (851.2m)
Underflow Threshold Bandwidth           : 10% (851.2m)
Overflow Threshold Activate Bandwidth   : 0
Underflow Threshold Activate Bandwidth  : 0
Overflow Limit                          : 1
Underflow Limit                         : 0
Max. Bandwidth Exceed Limit             : 1
-------------------------------------------------------------------------------
Max-BW-exceed-limit action              : notify
Resignal-failure-action                 : notify
Monitor Bandwidth                       : No
-------------------------------------------------------------------------------
Minimum Average Bandwidth               : 6.5g
Maximum Average Bandwidth               : 0
```

```
Total Overflow Count                    : 0
Consecutive Overflow Count              : 0
Consecutive Eligible Overflow Count     : 0
Total Underflow Count                   : 5
Consecutive Underflow Count             : 5
Consecutive Eligible Underflow Count    : 5
Max. Bandwidth Exceed Count             : 0
Teardown Count                          : 0
------------------------------------------------------------------------------
Last Bandwidth                          : 6.2g
Last Requested Bandwidth                : 8.5g
Last Signaled Bandwidth                 : 8.5g
Current Bandwidth                       : 8.5g
Highest Bandwidth                       : 8.6g
------------------------------------------------------------------------------
Time for Next Sample request            : 17 seconds
Time for Next Adjustment                : 0 seconds
Time of Last Bandwidth Request          : 2024 Jun 25 11:30:42
Time of Last Bandwidth Signal            : 2024 Jun 25 11:30:42
Time of Last Adjustment                 : 2024 Jun 25 11:30:42
Time of Highest Bandwidth Marked        :2024 Jun 25 11:24:35
------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments        : 11
Successful Adjustments                  : 11
Failed Adjustments                      : 0
------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
  [Sample 1-5]       : 6.5g     6.3g     6.4g     6.4g     6.4g

#show rsvp trunk auto-bandwidth
--------------------+--------+--------+--------+----------+----------+----------+------
----+------------+--------------------
Trunk-Name          Trunk    LSP      Last     Requested  Signaled   Current    Highest
Adjust-Time  Last-Adjust
                    ID       ID       BW       BW         BW         BW         BW
Left(sec)    Time
--------------------+--------+--------+--------+----------+----------+----------+------
----+------------+--------------------
PE1-PE2                  5002    2202    8.5g    6.5g      6.5g      6.5g      8.6g
277          2024 Jun 25 11:30:42
```

Configure RSVP Auto Bandwidth on Boot on PE1 Router

1. Create auto-bandwidth Profile.

```
(config)#router rsvp
```

2. Configure Auto bandwidth on boot and set the values for sample interval, Adjust interval and Adjust interval count.

```
(config-router)#auto-bandwidth-on-boot 1 5 1
(config-router)#commit
```

Validation

Verify auto bandwidth on boot adjustments information as below:

```
#show running-config rsvp
!
router rsvp
 auto-bandwidth-on-boot 1 5 1
!
!
!
!
!
!
!
#
#show rsvp trunk auto-bandwidth

*** On boot auto bandwidth is in progress for 2 minutes, 3 seconds ***
--------------------+--------+--------+--------+----------+----------+----------+------
----+------------+--------------------
Trunk-Name          Trunk    LSP      Last     Requested  Signaled   Current    Highest
Adjust-Time  Last-Adjust
                    ID       ID       BW       BW         BW         BW         BW
Left(sec)    Time
--------------------+--------+--------+--------+----------+----------+----------+------
----+------------+--------------------
PE-1_to_PE-2_1     5002     2202     0        0          0          200m       144.6m
NA           NA
```

Commands for RSVP Auto-Bandwidth

The RSVP auto-bandwidth uses the following configuration commands.

rsvp-auto-bandwidth

Use this command to configure an auto bandwidth profile. The profile will have default settings if any parameter not configured explicitly. User can configure parameters to their need within auto bandwidth profile.

Use `no` parameter of this command to delete auto bandwidth profile.

Command Syntax

```
rsvp-auto-bandwidth PROFILENAME
no rsvp-auto-bandwidth PROFILENAME
```

Parameters

`<PROFILE_NAME>`	Specifies the name assigned to the auto-bandwidth profile during configuration. The profile name can be a maximum of 64 characters in length.

Default

None

Command Mode

Config mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure an auto-bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#commit
(config-auto-bandwidth)#exit
(config)#
```

The following example describes how to delete the auto bandwidth profile:

```
#configure terminal
(config)#no rsvp-auto-bandwidth bwp
(config)#commit
```

sample-interval

Use this command to configure a sample interval value in minutes on the auto bandwidth profile. Sample interval determines the frequency at which rate samples collected from associated trunks. Sample interval must not be configured more than adjust interval as no samples can be collected within an adjustment cycle in such case.

Note: Sample interval timers run per auto bandwidth profile and not per associated trunks. So, in case of bandwidth adjustments on trunks before adjustment cycle completion will leave the newly formed session with less number of samples in the remaining part of adjustment cycle. In order to avoid very few samples being processed, minimum-samples command shall be configured in absolute or percentage format.

Use the `no` parameter to remove the sample interval configuration.

Command Syntax

```
sample-interval <1 - 10080>
no sample-interval
```

Parameters

`<1-10080>`	Specifies the sample interval value in minutes.

Default

5 minutes

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure sample interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#sample-interval 2
(config-auto-bandwidth)#commit
```

The following example describes how to remove configured sample interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no sample-interval
(config-auto-bandwidth)#commit
```

adjust-interval

Use this command to configure a adjust interval value in minutes on the auto bandwidth profile. Adjust interval determines the duration of the adjustment cycle. Bandwidth update decisions for active session of associated trunks are taken after every adjustment cycle. Adjust interval must not be configured less than sample interval as no samples can be collected within an adjustment cycle in such case.

Note: Adjust interval timers run per auto bandwidth profile and not per associated trunks. So, in case of bandwidth adjustments on trunks before adjustment cycle completion will leave the newly formed session with less number of samples in the remaining part of adjustment cycle. In order to avoid very few samples being processed, minimum-samples command shall be configured in absolute or percentage format.

Use the `no` parameter to remove the adjust interval configuration.

Command Syntax

```
adjust-interval <5 - 10080>
no adjust-interval
```

Parameters

`<5-10080>`	Specifies the adjust interval value in minutes.

Default

30 minutes

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure adjust interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#adjust-interval 60
(config-auto-bandwidth)#commit
```

The following example describes how to remove configured adjust interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no adjust-interval
(config-auto-bandwidth)#commit
```

minimum-bandwidth

Use this command to configure minimum bandwidth on the auto bandwidth profile. Even when traffic flow is much lesser than minimum bandwidth, LSP will be reserved with the configured minimum bandwidth during bandwidth adjustment process. When auto bandwidth profile associated with trunk, LSP will be signaled with minimum bandwidth when initial bandwidth is not configured in the profile.

Use the `no` parameter to remove the minimum bandwidth configuration from the profile.

Command Syntax

```
minimum-bandwidth BANDWIDTH
```

```
no minimum-bandwidth
```

Parameters

`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure minimum bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#minimum-bandwidth 100m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the minimum bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no minimum-bandwidth
(config-auto-bandwidth)#commit
```

maximum-bandwidth

Use this command to configure maximum bandwidth on the auto bandwidth profile. Even when traffic flow is much higher than maximum bandwidth, LSP will be reserved with the configured maximum bandwidth during bandwidth adjustment process. Operator notification is generated if the traffic rate samples collected are higher than the maximum bandwidth but the reservation is limited to maximum bandwidth.

Note: When maximum bandwidth is configured, even a single traffic rate sample crossing the maximum bandwidth will trigger an MBB with maximum bandwidth reserved. If user doesn't wish to trigger an MBB for single sample of maximum bandwidth exceed, maximum-bandwidth-exceed-limit shall be configured with a value to mention the number of consecutive samples to cross maximum bandwidth to take further action.

Use the `no` parameter to remove the maximum bandwidth configuration from the profile.

Command Syntax

```
maximum-bandwidth BANDWIDTH
no maximum-bandwidth
```

Parameters

`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure maximum bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#maximum-bandwidth 900m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the maximum bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no maximum-bandwidth
(config-auto-bandwidth)#commit
```

initial-bandwidth

Use this command to configure initial bandwidth on the auto bandwidth profile. When auto bandwidth profile associated with trunk, LSP will be signalled with initial bandwidth when initial bandwidth is configured in the profile. For trunks which are already associated with auto bandwidth profile and the system going through reload, initial bandwidth will not be applicable as on boot computation will trigger to update active sessions with bandwidth as per the on boot period traffic rate sample computation.

Use the `no` parameter to remove the initial bandwidth configuration from the profile.

Command Syntax

```
initial-bandwidth BANDWIDTH
no initial-bandwidth
```

Parameters

`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure initial bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#initial-bandwidth 500m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the initial bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no initial-bandwidth
(config-auto-bandwidth)#commit
```

underflow-threshold

Use this command to configure underflow threshold in percentage or absolute value format on the auto bandwidth profile. Underflow threshold sets the amount of reduction in traffic rate sample to detect an eligible underflow. As an example, absolute underflow threshold 10m when current bandwidth is 200m means, a traffic rate sample of 190.1m will not be considered eligible underflow sample and a sample of 189.9m will be considered eligible underflow sample.

When all the traffic rate samples collected for an auto bandwidth profile associated trunk cross underflow threshold in an adjustment cycle, then the highest eligible traffic rate sample will be considered to re-signal the session with new bandwidth at the end of an adjustment cycle.

Constraints like underflow-limit and underflow-threshold-activate-bandwidth will add additional logic on how bandwidth update action is taken. This will be discussed in respective sections.

If underflow threshold is not configured, then minor reduction in traffic rate sample also will be considered as eligible underflow bandwidth sample. So, underflow and overflow threshold is a recommended configuration even though it is not mandatory.

When underflow threshold is configured in percentage, the threshold will be computed based on the current bandwidth and the percentage value. Example, underflow threshold 10% for a current bandwidth of 100m means a sample of 90m

or lesser will be considered eligible underflow sample. Underflow threshold can be configured either as absolute value or in percentage but not both.

Use the `no` parameter to remove the underflow bandwidth configuration from the profile.

Command Syntax

```
underflow-threshold (percent <1-100>) | (absolute BANDWIDTH)
no underflow-threshold (percent | absolute)
```

Parameters

`<1-100>`	Specifies the underflow threshold value in percentage.
`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow bandwidth in percentage format in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#underflow-threshold percent 10
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no underflow-threshold percent
(config-auto-bandwidth)#commit
```

overflow-threshold

Use this command to configure overflow threshold in percentage or absolute value format on the auto bandwidth profile. Overflow threshold sets the amount of increase in traffic rate sample required to detect an eligible overflow. As an example, absolute overflow threshold 10m when current bandwidth is 200m means, a traffic rate sample of 209.9m will not be considered eligible overflow sample and a sample of 210.1m will be considered eligible overflow sample.

When a traffic rate sample collected for an auto bandwidth profile associated trunk crosses overflow threshold in an adjustment cycle, then the highest eligible traffic rate sample will be considered to re-signal the session with new bandwidth at the end of adjustment cycle.

Constraints like overflow-limit and overflow-threshold-activate-bandwidth will add additional logic on how bandwidth update action is taken. This will be discussed in respective sections.

If overflow threshold is not configured, then minor increase in traffic rate sample also will be considered as eligible overflow bandwidth sample. So, underflow and overflow threshold is a recommended configuration even though it is not mandatory.

When overflow threshold is configured in percentage, the threshold will be computed based on the current bandwidth and the percentage value. Example, overflow threshold 10% for a current bandwidth of 100m means a sample of 110m or more will be considered eligible overflow sample. Overflow threshold can be configured either as absolute value or in percentage but not both.

Use the `no` parameter to remove the underflow bandwidth configuration from the profile.

Command Syntax

```
overflow-threshold (percent <1-100>) | (absolute BANDWIDTH)
no overflow-threshold (percent | absolute)
```

Parameters

`<1-100>`	Specifies the overflow threshold value in percentage.
`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow bandwidth in absolute format in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#overflow-threshold absolute 10m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no overflow-threshold absolute
(config-auto-bandwidth)#commit
```

underflow-threshold-activate-bandwidth

Use this command to configure absolute bandwidth range to allow bandwidth re-signalling when underflow threshold and underflow limit criteria matched. This configuration helps to limit the underflow bandwidth reservation update for certain range of bandwidth.

As an example, if the current bandwidth is 500m and the underflow threshold is 10%. So, normally, if all traffic rate samples collected are in the range of 400m to 450m, session will be re-signalled to reserve new bandwidth. However, if underflow-threshold-activate-bandwidth is configured as 300m, then the traffic rate samples in the range of 400m to

450m will not trigger bandwidth update. Only when the traffic rate samples are less than 300m, then it will be considered as eligible sample.

The configuration creates an absolute bandwidth range for underflow samples to be eligible. The bandwidth range for underflow eligibility will be minimum bandwidth (or zero when minimum bandwidth is not configured) to underflow-threshold-activate-bandwidth value. When this command is not configured, there won't be any such absolute range and only underflow-threshold and underflow-limit will be considered for computation, if configured.

Use the `no` parameter to remove the underflow threshold activate bandwidth configuration from the profile.

Command Syntax

```
underflow-threshold-activate-bandwidth BANDWIDTH
no underflow-threshold-activate-bandwidth
```

Parameters

`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow threshold activate bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#underflow-threshold-activate-bandwidth 500m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow threshold activate bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no underflow-threshold-activate-bandwidth
(config-auto-bandwidth)#commit
```

overflow-threshold-activate-bandwidth

Use this command to configure absolute bandwidth range to allow bandwidth re-signalling when overflow threshold and overflow limit criteria matched. This configuration helps to limit the overflow bandwidth reservation update for certain range of bandwidth.

As an example, if the current bandwidth is 100m and the overflow threshold is 10%. Normally, if a traffic rate sample collected is in the range of 110m to 150m, session will be re-signalled to reserve new bandwidth. However, if overflow-threshold-activate-bandwidth is configured as 300m, then the traffic rate samples in the range of 110m to 150m will not trigger bandwidth update. Only when the traffic rate samples are more than 300m, then it will be considered as eligible sample.

The configuration creates an absolute bandwidth range for overflow samples to be eligible. The bandwidth range for overflow eligibility will be overflow-threshold-activate-bandwidth value to a practical infinity. When this command is not configured, there won't be any such absolute range and only overflow-threshold and overflow-limit will be considered for computation, if configured.

Use the `no` parameter to remove the overflow threshold activate bandwidth configuration from the profile.

Command Syntax

```
overflow-threshold-activate-bandwidth BANDWIDTH
no overflow-threshold-activate-bandwidth
```

Parameters

`BANDWIDTH`	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure overflow threshold activate bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#overflow-threshold-activate-bandwidth 500m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the overflow threshold activate bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no overflow-threshold-activate-bandwidth
(config-auto-bandwidth)#commit
```

underflow-limit

Use this command to configure underflow limit on the auto bandwidth profile. When underflow limit is configured, if the traffic rate samples collected on the associated session consecutively crosses underflow threshold for underflow limit times, then the bandwidth adjustment will be triggered immediately without waiting for adjustment cycle completion. When underflow-threshold-activate-bandwidth is configured, even this criteria is considered to mark a sample as eligible underflow sample.

Only when underflow limit is configured, underflow adjustment may happen before the completion of adjustment cycle. Otherwise, underflow adjustment considered only at the completion of adjustment cycle when all samples found to be eligible underflow sample.

Use the `no` parameter to remove the underflow limit configuration from the profile.

Command Syntax

```
underflow-limit <1-10080>
no underflow-limit
```

Parameters

<1-10080>	Specifies the underflow limit value for consecutive eligible samples.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow limit in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#underflow-limit 3
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow limit configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no underflow-limit
(config-auto-bandwidth)#commit
```

overflow-limit

Use this command to configure overflow limit on the auto bandwidth profile. When overflow limit is configured, if the traffic rate samples collected on the associated session consecutively crosses overflow threshold for overflow limit times, then the bandwidth adjustment will be triggered immediately without waiting for adjustment cycle completion. When overflow-threshold-activate-bandwidth is configured, even this criteria is considered to mark a sample as eligible underflow sample.

Only when overflow limit is configured, overflow adjustment may happen before the completion of adjustment cycle. Otherwise, overflow adjustment considered only at the completion of adjustment cycle when a sample found to be eligible overflow sample.

If the traffic rate sample crosses maximum bandwidth, then maximum-bandwidth-exceed-limit configuration comes into picture and by default, a single sample crossing maximum bandwidth triggers bandwidth update. This situation is different from overflow scenario.

Use the `no` parameter to remove the overflow limit configuration from the profile.

Command Syntax

```
overflow-limit <1-10080>
no overflow-limit
```

Parameters

<1-10080>	Specifies the overflow limit value for consecutive eligible samples.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure overflow limit in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#overflow-limit 3
(config-auto-bandwidth)#commit
```

The following example describes how to remove the overflow limit configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no overflow-limit
(config-auto-bandwidth)#commit
```

maximum-bandwidth-exceed-limit

Use this command to configure maximum bandwidth exceed limit on the auto bandwidth profile. When maximum bandwidth exceed limit is configured, if the traffic rate samples collected on the associated session consecutively crosses maximum bandwidth for maximum-bandwidth-exceed-limit times, then the action will be triggered immediately without waiting for adjustment cycle completion. When maximum-bandwidth-exceed-limit is not configured, a single sample exceeding maximum bandwidth will trigger an action which is re-signal with updated bandwidth or restart the session with initial or minimum bandwidth based on the action configured.

When maximum bandwidth is not configured, maximum bandwidth exceed limit configuration doesn't have any significance. Overflow limit and maximum bandwidth exceed limits are independent commands with different significance with latter associated with maximum bandwidth.

Use the `no` parameter to remove the maximum bandwidth exceed limit configuration from the profile.

Command Syntax

```
maximum-bandwidth-exceed-limit <1-10080>
no maximum-bandwidth-exceed-limit
```

Parameters

<1-10080>	Specifies the maximum bandwidth exceed limit value for consecutive eligible samples.

Default

1

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure maximum bandwidth exceed limit in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#maximum-bandwidth-exceed-limit 2
(config-auto-bandwidth)#commit
```

The following example describes how to remove the maximum bandwidth exceed limit configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no maximum-bandwidth-exceed-limit
(config-auto-bandwidth)#commit
```

maximum-bandwidth-exceed-action

Use this command to configure maximum bandwidth exceed action on the auto bandwidth profile. When the traffic rate samples collected on the associated session consecutively crosses maximum bandwidth for maximum-bandwidth-exceed-limit times (or one time if limit is not configured), then the action to be triggered will be decided by this configuration. If not configured, default action is to re-signal the session with maximum bandwidth or ignore if session is already signalled with maximum bandwidth. In any case, user will be notified about the maximum bandwidth being exceeded. However, with exceed action configured as teardown, session will be released and restarted with initial bandwidth or minimum bandwidth if initial bandwidth is not configured.

This action will lead to service interruption if there are no alternate transport. So, this configuration is recommended to be used with full awareness of the impact.

Use the `no` parameter to remove the maximum bandwidth exceed action configuration from the profile.

Command Syntax

```
maximum-bandwidth-exceed-action (teardown)
no maximum-bandwidth-exceed-action
```

Parameters

`teardown`	Teardown the session exceeding maximum bandwidth.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure maximum bandwidth exceed action in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#maximum-bandwidth-exceed-action teardown
(config-auto-bandwidth)#commit
```

The following example describes how to remove the maximum bandwidth exceed action configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no maximum-bandwidth-exceed-action
(config-auto-bandwidth)#commit
```

resignal-failure-action

Use this command to configure an action on the auto bandwidth profile when the bandwidth update re-signalling fails on the associated session. By default, if re-signalling fails (3 attempts) for the updated bandwidth, it will be noted down as re-signalling failure and session will continue with its current bandwidth reservation. If severe actions to be taken on such re-signal failure, then teardown action can be configured which will release the current session and restart freshly with initial bandwidth or minimum bandwidth when initial bandwidth is not configured.

This action will lead to service interruption if there are no alternate transport. So, this configuration is recommended to be used with full awareness of the impact.

Use the `no` parameter to remove the re-signal failure action configuration from the profile.

Command Syntax

```
resignal-failure-action (teardown)
```

```
no resignal-failure-action
```

Parameters

`teardown`	Specifies the teardown the session when re-signalling with new bandwidth fails.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure re-signal failure action in a profile:

```
#configure terminal
```

```
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#resignal-failure-action teardown
(config-auto-bandwidth)#commit
```

The following example describes how to remove the re-signal failure action configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no resignal-failure-action
(config-auto-bandwidth)#commit
```

sync-bandwidth

Use this command to configure bandwidth synchronization for primary and secondary sessions of an auto bandwidth profile associated trunk. With this configuration, in case the associated trunk is configured with primary and secondary sessions, every time primary session goes through a bandwidth update, secondary session also will be re-signalled with primary session's bandwidth. Thus, secondary path is determined with proper reservation constraints to ensure it is in the correct bandwidth reserved state when traffic switches to secondary.

Use the `no` parameter to remove synchronise bandwidth configuration from the profile.

Command Syntax

```
sync-bandwidth
```

```
no sync-bandwidth
```

Parameters

None

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure synchronize bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#sync-bandwidth
(config-auto-bandwidth)#commit
```

The following example describes how to remove the synchronise bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no sync-bandwidth
(config-auto-bandwidth)#commit
```

monitor-bandwidth

Use this command to configure only monitor the traffic rate samples and computation without taking any action. This command can be used to monitor the traffic behaviour without updating the active sessions. With this configuration, in case of overflow, underflow, adjustment cycle completion time computation results, maximum bandwidth exceed, etc., notification is provided without taking any action.

Use the `no` parameter to remove monitor bandwidth configuration from the profile.

Command Syntax

```
monitor-bandwidth
no monitor-bandwidth
```

Parameters

None

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure monitor bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#monitor-bandwidth
(config-auto-bandwidth)#commit
```

The following example describes how to remove the monitor bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no monitor-bandwidth
(config-auto-bandwidth)#commit
```

minimum-samples

Use this command to configure the minimum samples required in an adjustment cycle for bandwidth processing. Sample timers and Adjust timers are executed per auto bandwidth profile and not per associated trunk. Thus, there are scenarios of a trunk going through a bandwidth update few minutes ago and again ends up with adjustment cycle completion processing with very few samples collected. In order to avoid such scenarios, minimum samples required in an adjustment cycle to process the bandwidth shall be configured.

Configuration is accepted in both absolute and in percentage format. This gives user the flexibility to choose the format that suites their need. If sample interval and adjust interval expected to be fixed, then absolute configuration helps providing the requirement of exact number of minimum samples required to process. If exact number isn't important and there are chances of changing adjust interval or sample interval in future, then percentage format can be chosen. However, only one of the formats can be configured.

By default, even if there is one traffic rate sample during adjustment cycle completion, bandwidth will be processed. So, it will be recommended to have this configuration if users are keen on minimum of certain samples to be considered for bandwidth computation.

Use the `no` parameter to remove the minimum sample configuration from the profile.

Command Syntax

```
minimum-samples (percent <1-100>) | (absolute <1-10080>)
no underflow-limit (percent | absolute)
```

Parameters

`<1-10080>`	Specifies the absolute value for minimum samples required in an adjustment cycle.
`<1-100>`	Specifies the minimum sample percentage required in an adjustment cycle.

Default

1 sample

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to configure minimum samples in percentage format in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#minimum-samples percent 70
(config-auto-bandwidth)#commit
```

The following example describes how to remove the minimum samples configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no minimum-samples percent
(config-auto-bandwidth)#commit
```

auto-bandwidth

Use this command to attach an auto bandwidth profile to a trunk. When the auto bandwidth profile is attached to the trunk, active session will be re-signalled with initial bandwidth configured in the auto bandwidth profile or minimum bandwidth configured if initial bandwidth is not configured. Bandwidth update will be triggered only if there is variation in the bandwidth to be initialized. Attaching or detaching an auto bandwidth profile doesn't trigger any session flap and doesn't cause traffic impact.

When an auto bandwidth profile is associated with first trunk, sample interval and adjust interval timers will start and are stopped when the profile is removed from the last trunk.

Manual bandwidth configuration for the sessions and auto bandwidth profile attach are mutually exclusive and the configuring both of them on a trunk is not allowed.

Use the `no` parameter to remove the auto bandwidth profile from the trunk.

Command Syntax

```
auto-bandwidth PROFILENAME
no auto-bandwidth PROFILENAME
```

Parameters

PROFILENAME	Specifies the name of the auto bandwidth profile.

Default

None

Command Mode

Trunk mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to associate an auto bandwidth profile to a trunk:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#exit
(config)#rsvp-trunk t1
(config-trunk)#auto-bandwidth bwp
(config-trunk)#commit
```

The following example describes how to remove the auto bandwidth profile from the trunk:

```
#configure terminal
(config)#rsvp-trunk t1
(config-trunk)#no auto-bandwidth bwp
(config-trunk)#commit
```

auto-bandwidth-on-boot

Use this command to configure on boot sample interval, adjust interval and number of adjustment cycles. When the system is reloaded and comes up, all active sessions of trunks associated with auto bandwidth profiles run a relatively faster adjustment cycle with quicker sample collection to settle the sessions with accurate bandwidth reservation.

By default, sample interval is 1 minute, adjust interval is 5 minutes and the adjustment cycle runs one time. After the adjustment cycle completion, samples of each associated trunks computed to re-signal the sessions with updated bandwidth. Then the auto bandwidth profile based adjustment cycle starts. If user wishes to run the boot up time rigorous sample computation for longer duration or multiple rounds, then it shall be configured. The configurations will apply from system reload if the configuration is saved. Properties of auto bandwidth profiles will not be applied during boot up time computation.

Use the `no` parameter to remove the auto bandwidth profile from the trunk.

Command Syntax

```
auto-bandwidth-on-boot <1-10080> <1-10080> <1-10>
no auto-bandwidth-on-boot
```

Parameters

<1-10080>	On boot sample interval value in minutes.
<1-10080>	On boot adjustment interval value in minutes.
<1-10>	Specifies the number of adjustment cycles to run on boot.

Default

Sample interval 1 minute, adjust interval 5 minutes and 1 adjustment cycle.

Command Mode

Router mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure on boot auto bandwidth parameters:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bandwidth-on-boot 1 10 3
(config-router)#commit
```

The following example describes how to reset on boot auto bandwidth parameters:

```
#configure terminal
(config)#router rsvp
(config-router)#no auto-bandwidth-on-boot
(config-router)#commit
```

force-auto-bandwidth-adjustment

Use this command to force a bandwidth adjustment on a trunk associated with auto bandwidth profile. When the command is executed without bandwidth value mentioned, traffic rate samples collected till the time are used to compute the bandwidth to be adjusted. In case of bandwidth value mentioned in the command, the bandwidth is verified for eligibility and bandwidth update will be triggered.

Command Syntax

```
rsvp-trunk TRUNKNAME force-auto-bandwidth-adjustment (BANDWIDTH|)
```

Parameters

TRUNKNAME	Specifies the name of the trunk to go through forced bandwidth adjustment.
BANDWIDTH	Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how you can force a bandwidth adjustment for a trunk with an auto bandwidth profile:

```
#rsvp-trunk t1 force-auto-bandwidth-adjustment
```

clear rsvp auto-bandwidth

Use this command to reset the auto bandwidth adjustment cycle by clearing all the traffic samples collected by the associated trunks and by restarting sample and adjust timers. If auto bandwidth profile name is not mentioned, then all trunks associated with any auto bandwidth profile will be reset and computation will start freshly.

Command Syntax

```
clear rsvp auto-bandwidth (PROFILENAME|)
```

Parameters

PROFILENAME	Specifies the name of the auto bandwidth profile.

Default

None

Command Mode

Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to restart processing of an auto bandwidth profile:

```
#clear rsvp auto-bandwidth bwp
```

clear rsvp trunk auto-bandwidth-statistics

Use this command to clear the statistics maintained on a trunk associated with auto bandwidth profile. Statistics will be mainly the highest watermarked bandwidth, last adjusted bandwidth, how many times adjustment triggered, status of the adjustment trigger, etc. This command will only clear the auto bandwidth statistics for the trunk and doesn't impact the operation of auto bandwidth including the traffic rate samples collected for the current adjustment cycle.

Command Syntax

```
clear rsvp trunk TRUNKNAME auto-bandwidth-statistics
```

Parameters

TRUNKNAME	Specifies the name of the trunk associated with auto bandwidth profile.

Default

None

Command Mode

Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to clear the auto bandwidth statistics on a trunk which is associated with the auto bandwidth profile:

```
#clear rsvp trunk t1 auto-bandwidth-statistics
```

Show Commands for RSVP

show rsvp auto-bandwidth

Use this command to display auto bandwidth profile specific information.

Command Syntax

```
show rsvp auto-bandwidth
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing all the auto bandwidth profiles:

```
 #show rsvp auto-bandwidth

Profile Name : bwp
------------------------------------------------------------------------------
Sample Interval                          : 5 minutes (due in 4 minutes)
Adjust Interval                          : 30 minutes (due in 29 minutes)
Minimum Samples required for processing  : 1
Initial Bandwidth                        : 0
Minimum bandwidth                        : 0
Maximum bandwidth                        : 100m
Underflow Threshold Bandwidth            : 5m
```

```
Overflow Threshold Bandwidth                : 5m
Underflow Threshold Activate Bandwidth      : 0
Overflow Threshold Activate Bandwidth       : 0
Overflow Limit                              : 3
Underflow Limit                             : 3
Maximum Bandwidth Exceed Limit              : 1
Maximum Bandwidth Exceed Action             : notify
Re-signal Failure Action                    : notify
Sync Bandwidth                              : No
Monitor Bandwidth                           : No
No. of trunks associated                    : 1
```

show rsvp auto-bandwidth detail

Use this command to display a specific auto bandwidth profile information or all auto bandwidth profile information along with associated trunk details.

Command Syntax

```
show rsvp auto-bandwidth (PROFILENAME | detail)
```

Parameters

`PROFILENAME`	Specifies the name of the auto bandwidth profile.
`detail`	Specifies detailed information of all the auto bandwidth profiles.

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example is for viewing all the auto bandwidth profiles along with associated trunks:

```
#show rsvp auto-bandwidth detail
```

```
Profile Name : bwp
------------------------------------------------------------------------------
Sample Interval                           : 5 minutes (due in 4 minutes)
Adjust Interval                           : 30 minutes (due in 29 minutes)
Minimum Samples required for processing   : 1
Initial Bandwidth                         : 0
Minimum bandwidth                         : 0
Maximum bandwidth                         : 100m
Underflow Threshold Bandwidth             : 5m
Overflow Threshold Bandwidth              : 5m
Underflow Threshold Activate Bandwidth    : 0
Overflow Threshold Activate Bandwidth     : 0
Overflow Limit                            : 3
Underflow Limit                           : 3
Maximum Bandwidth Exceed Limit            : 1
Maximum Bandwidth Exceed Action           : notify
Re-signal Failure Action                  : notify
Sync Bandwidth                            : No
Monitor Bandwidth                         : No
```

```
No. of trunks associated                  : 1
------------+--------+--------+--------+----------+----------+----------+----------+-----------
Trunk-Name    Trunk    LSP      Last     Requested  Signaled   Current    Highest    LastAdjust
            ID       ID       BW       BW         BW         BW         BW         Time                -
------------+--------+--------+--------+----------+----------+----------+----------+-----------
t1            5001     2201     10.1m    22.5m      22.5m      22.5m      35.5m      2024 Jul 23
```

show rsvp trunk auto-bandwidth

Use this command to display the information of all the trunks associated with the auto bandwidth profile. This show command will display high level information like what is the last bandwidth, current bandwidth, last adjustment time, time left in adjustment cycle in seconds, etc.

Command Syntax

```
show rsvp trunk auto-bandwidth
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing an auto bandwidth summary of all the trunks associated with auto bandwidth profile:

```
      #show rsvp trunk auto-bandwidth
-----------+------+-----+------+----------+---------+--------+--------+------------+-----------
Trunk-Name  Trunk  LSP   Last   Requested  Signaled  Current  Highest  Adjust-Time  Last-Adjust
           ID     ID    BW     BW         BW        BW       BW       Left(sec)    Time
-----------+------+-----+------+----------+---------+--------+--------+------------+-----------
t1           5001   2201  10.1m  22.5m      22.5m     22.5m    35.5m    1142         2024 Jul 23
```

show rsvp trunk auto-bandwidth detail

Use this command to display the information of a trunk or all the trunks associated with the auto bandwidth profile. This command will provide detailed information of the auto bandwidth related statistics on the trunk as well as details of traffic rate samples collected in an adjust cycle and the time left for next sample collection, etc.

Command Syntax

```
show rsvp trunk auto-bandwidth (TRUNKNAME | detail)
```

Parameters

TRUNKNAME	Specifies the name of the particular trunk to display auto-bandwidth details for.

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing the auto bandwidth details of all the trunks associated with auto bandwidth profile

```
#show rsvp trunk auto-bandwidth detail
:Session: t1-Primary, Tunnel-id: 5001, LSP-ID: 2201, Egress: 2.2.2.2
------------------------------------------------------------------------------
Sample Interval                          : 5 minutes
Adjustment Interval                      : 30 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth                 : 0
Minimum Bandwidth                        : 0
Maximum Bandwidth                        : 100m
Overflow Threshold Bandwidth             : 5m
Underflow Threshold Bandwidth            : 5m
Overflow Threshold Activate Bandwidth    : 0
Underflow Threshold Activate Bandwidth   : 0
Overflow Limit                           : 3
Underflow Limit                          : 3
Max. Bandwidth Exceed Limit              : N/A
------------------------------------------------------------------------------
Max-BW-exceed-limit action               : notify
Resignal-failure-action                  : notify
Monitor Bandwidth                        : No
------------------------------------------------------------------------------
Minimum Average Bandwidth                : 0
Maximum Average Bandwidth                : 22.5m
Total Overflow Count                     : 1
Consecutive Overflow Count               : 1
Consecutive Eligible Overflow Count      : 1
Total Underflow Count                    : 0
Consecutive Underflow Count              : 0
Consecutive Eligible Underflow Count     : 0
Max. Bandwidth Exceed Count              : 0
Teardown Count                           : 0
------------------------------------------------------------------------------
Last Bandwidth                           : 10.2m
Last Requested Bandwidth                 : 15.6m
Last Signaled Bandwidth                  : 15.6m
Current Bandwidth                        : 15.6m
Highest Bandwidth                        : 35.3m
------------------------------------------------------------------------------
Time for Next Sample request             : 1 minutes, 20 seconds
Time for Next Adjustment                 : 16 minutes, 30 seconds
Time of Last Bandwidth Request           : 2024 Jul 23 11:32:44
Time of Last Bandwidth Signal            : 2024 Jul 23 11:32:44
Time of Last Adjustment                  : 2024 Jul 23 11:32:44
Time of Highest Bandwidth Marked         : 2024 Jul 23 11:14:37
------------------------------------------------------------------------------
Total Auto-Bandwidth Adjustments         : 2
```

```
Successful Adjustments                  : 2
Failed Adjustments                      : 0
--------------------------------------------------------------------------------
Samples collected in the current adjustment cycle:
  [Samples 1-5]      : 17.5m      18.3m      22.5m
```

CHAPTER 6 Y.1731 and CFM Over EVPN ELINE Single Home

Overview

The Single Home EVPN ELINE Y.1731 CFM over Sub-interface feature enables the monitoring and management of Ethernet Virtual Private Network (EVPN) E-Line services using the Y.1731 Connectivity Fault Management (CFM) protocol over sub-interfaces. This feature enhances fault detection and performance monitoring capabilities for EVPN E-Line services, allowing network operators to ensure high availability and reliability of their networks. By extending Y.1731 CFM functionality to sub-interfaces in single home EVPN E-Line deployments, this feature provides comprehensive end-to-end visibility and control, enabling proactive fault detection, isolation, and troubleshooting.

Feature Characteristics

- Utilizes sub-interfaces to partition Ethernet traffic within the Single Home EVPN ELINE architecture, enabling efficient service delivery and management.
- Implements EVPN ELINE architecture with single-homing capabilities, facilitating the creation of Ethernet Virtual Private Networks with simplified configurations and reduced complexity.
- Provides robust fault detection mechanisms to identify connectivity issues, link failures, and service disruptions in Ethernet networks.

Benefits

- Provides detailed insights into Ethernet service performance, enabling proactive monitoring and optimization of network resources.
- Minimizes service downtime by promptly detecting and resolving faults, ensuring uninterrupted service delivery and customer satisfaction.
- Optimizes network resource utilization and bandwidth allocation by identifying and addressing connectivity issues in a timely manner.
- Facilitates rapid fault identification and isolation, accelerating troubleshooting processes and reducing mean time to repair (MTTR).
- Ensures compliance with Service Level Agreements (SLAs) by maintaining service quality metrics within defined thresholds and objectives.

Prerequisites

Ensure that the network devices (routers, switches) support Y.1731 CFM functionality and Single Home EVPN ELINE configuration.

Verify that the devices are running compatible software versions that include support for these features.

Configuration

Configure Single Home EVPN ELINE Y.1731 CFM over Sub-interface for enhanced fault management in EVPN networks.

Topology

The topology consists of two Customer Edge devices (CE1 and CE2) connected to Provider Edge devices (PE1 and PE2) through sub-interfaces. The Provider Edge devices are interconnected through Provider devices (P1 and P2). Y.1731 functionality is implemented over these sub-interfaces, allowing for fault detection and performance monitoring of Ethernet connectivity between the customer sites.

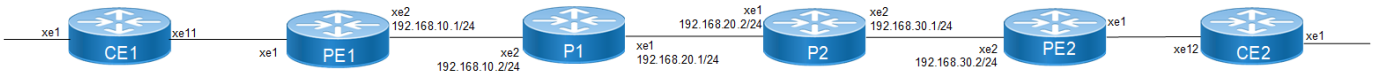

Figure 6-4: EVPN ELINE Over Sub-interface-Single Home

Perform the following configurations to configure Single Home EVPN ELINE Y.1731 CFM over Sub-interface:

1. On Customer Edge (CE) Nodes (CE1 and CE2), configure the interface xe1 and set it as a switchport with a load interval of (`30 seconds`):

```
CE1(config)#interface xe1
CE1(config-if)#switchport
CE1(config-if)#load-interval 30
CE1(config-if)#commit
CE1(config-if)#exit
```

Note: Similarly follow the same steps to configure xe11(CE1) and xe12(CE2).

2. Create sub-interface (`xe1.2001`) adding the VLAN:

```
CE1(config)#interface xe1.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit

CE1(config)#interface xe11.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit
```

3. Set up a cross-connect named (`test100`), specifying in and out interfaces:

```
CE1(config)#cross-connect test100
CE1(config-xc)#interface xe1.2001
CE1(config-xc)#interface xe11.2001
CE1(config-xc)#commit
```

4. Perform the following on PE1:

    1. Configure CFM related hardware profiles:

```
PE1(config)# hardware-profile filter cfm-domain-name-str enable
PE1(config)# hardware-profile statistics cfm-lm enable
PE1(config)# hardware-profile statistics cfm-ccm enable
PE1(config)#hardware-profile statistics cfm-slm enable
```

2. Configure the loopback interface with a secondary IP address(`1.1.1.1/32`):

```
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#commit
PE1(config-if)#exit
```

3. Configure LDP targeted peers:

```
PE1(config)#router ldp
PE1(config-router)#targeted-peer ipv4 4.4.4.4
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#commit
PE1(config-router)exit
```

4. Configure interface `xe2` with an IP address (`192.168.10.1/24`) and enable LDP:

```
PE1(config)#interface xe2
PE1(config-if)#load-interval 30
PE1(config-if)#ip address 192.168.10.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#commit
PE1(config-if)#exit
```

5. Configure OSPF routing, specify the OSPF router ID as (`1.1.1.1`), enable BFD on all interfaces, define the network (`1.1.1.1/32`) in area (`0.0.0.0`), and define the network (`192.168.10.0/24`) in area (`0.0.0.0`):

```
PE1(config)#router ospf 1
PE1(config-router)#ospf router-id 1.1.1.1
PE1(config-router)#bfd all-interfaces
PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0
PE1(config-router)#network 192.168.10.0/24 area 0.0.0.0
PE1(config-router)#commit
PE1(config-router)#exit
```

6. Enable EVPN MPLS globally and configure VTEP IP:

```
PE1(config)# evpn mpls enable
PE1(config)# commit
PE1(config)# evpn mpls vtep-ip-global 1.1.1.1
PE1(config)# commit
```

7. Configure BGP with the remote PE devices and activate EVPN:

```
PE1(config)# router bgp 100
PE1(config-router)# neighbor 4.4.4.4 remote-as 100
PE1(config-router)# neighbor 4.4.4.4 update-source lo
PE1(config-router)# address-family l2vpn evpn
PE1(config-router-af)# neighbor 4.4.4.4 activate
PE1(config-router-af)# exit
PE1(config-router)# exit
PE1(config)# commit
```

8. Configure MAC VRF with the appropriate RD and RT:

```
PE1(config)# mac vrf vrf2
PE1(config-vrf)# rd 1.1.1.1:2
PE1(config-vrf)# route-target both 2:2
PE1(config-vrf)# exit
```

9. Map the EVPN instance and VRF, specifying the EVPN ID:

```
PE1(config)# evpn mpls id 2 xconnect target-mpls-id 52
PE1(config-evpn-mpls)# host-reachability-protocol evpn-bgp vrf2
```

```
PE1(config-evpn-mpls)# evi-name test2
PE1(config-evpn-mpls)# commit
PE1(config-router-af)# exit
```

10. Configure access ports on PE1:

```
PE1(config)# interface xe1.2001 switchport
PE1(config-if)# encapsulation dot1q 2028
PE1(config-if)# access-if-evpn
PE1(config-acc-if-evpn)# map vpn-id 2
PE1(config-acc-if-evpn)# commit
```

11. Configure CFM MEP on PE1, define the FCM domain (`12346`), create MA, configure MEP, and configure Remote MEP Auto-discovery,set CC Interval `10ms`:

```
PE1(config)# ethernet cfm domain-type character-string domain-name12346
level 7 mip-creation default
PE1(config-ether-cfm)# service ma-type string ma-name 124
PE1(config-ether-cfm-ma)# ethernet cfm mep up mpid 20 active true
xe1.2001 vlan 2028
PE1(config-ether-cfm-ma-mep)# cc multicast state enable
PE1(config-ether-cfm-ma-mep)# exit-ether-ma-mep-mode
PE1(config-ether-cfm-ma)# rmep auto-discovery enable
PE1(config-ether-cfm-ma)# cc interval 10ms
PE1(config-ether-cfm-ma)# exit-ether-ma-mode
PE1(config-ether-cfm)# commit
```

12. Provide CFM configuration, define a delay measurement profile named DM, set the measurement interval to `1 second`, specify the number of intervals stored as `2`, configure the message period as `1 second`, define a loss measurement profile named LM, set the measurement type to LMM, set the measurement interval to `1 second`, specify the number of intervals stored as `3`, define a service level measurement profile named SLM, set the measurement type to SLM:

```
PE1(config)# ethernet cfm delay-measurement profile-name DM
PE1(config-cfm-dm)# measurement-interval 1
PE1(config-cfm-dm)# intervals-stored 2
PE1(config-cfm-dm)# message-period 1s
PE1(config-cfm-dm)# commit

PE1(config)# ethernet cfm loss-measurement profile-name LM
PE1(config-cfm-lm)# measurement-type lmm
PE1(config-cfm-lm)# measurement-interval 1
PE1(config-cfm-lm)# intervals-stored 3
PE1(config-cfm-lm)# message-period 1s
PE1(config-cfm-lm)# commit

PE1(config)# ethernet cfm loss-measurement profile-name SLM
PE1(config-cfm-lm)# measurement-type slm
PE1(config-cfm-lm)# measurement-interval 1
PE1(config-cfm-lm)# intervals-stored 3
PE1(config-cfm-lm)# message-period 1s
PE1(config-cfm-lm)# commit
```

Configuration Snapshot:

CE1:

```
interface xe1
switchport
load-interval 30
!
```

```
interface xe1.2001 switchport
encapsulation dot1q 2028
!

interface xe11.2001 switchport
encapsulation dot1q 2028
!
cross-connect test100
interface xe1.2001
interface xe11.2001
```

CE2:

```
interface xe1
switchport
load-interval 30
!
interface xe1.2001 switchport
encapsulation dot1q 2028
!
interface xe12.2001 switchport
encapsulation dot1q 2028
!
cross-connect test100
interface xe1.2001
interface xe12.2001
```

PE1:

```
interface lo
ip address 1.1.1.1/32 secondary
!
router ldp
targeted-peer ipv4 4.4.4.4
exit-targeted-peer-mode
!
interface xe2
load-interval 30
ip address 192.168.10.1/24
label-switching
enable-ldp ipv4
!
router ospf 1
ospf router-id 1.1.1.1
bfd all-interfaces
network 1.1.1.1/32 area 0.0.0.0
network 192.168.10.0/24 area 0.0.0.0
!
evpn mpls enable
evpn mpls vtep-ip-global 1.1.1.1
!
router bgp 100
neighbor 4.4.4.4 remote-as 100
neighbor 4.4.4.4 update-source lo
address-family l2vpn evpn
neighbor 4.4.4.4 activate
exit
!
```

```
mac vrf vrf2
rd 1.1.1.1:2
route-target both 2:2
!
evpn mpls id 2
xconnect target-mpls-id 52
host-reachability-protocol evpn-bgp vrf2
evi-name test2
!
interface xe1
switchport
load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-evpn
  map vpn-id 2
ethernet cfm domain-type character-string domain-name 12346 level 7
mipcreation none
  service ma-type string ma-name 124
    ethernet cfm mep up mpid 20 active true xe1.2001 vlan 2028
      cc multicast state enable
      exit-ether-ma-mep-mode
    rmep auto-discovery enable
    cc interval 10ms
    exit-ether-ma-mode
ethernet cfm loss-measurement profile-name SLM
measurement-type slm
measurement-interval 1
intervals-stored 3
message-period 1s
!
ethernet cfm loss-measurement profile-name LM
measurement-type lmm
measurement-interval 1
intervals-stored 3
message-period 1s
!
ethernet cfm delay-measurement profile-name DM
measurement-interval 1
intervals-stored 2
message-period 1s
```

PE2:

```
interface lo
 ip address 4.4.4.4/32 secondary

router ldp
 targeted-peer ipv4 1.1.1.1

interface xe2
 load-interval 30
 ip address 192.168.30.2/24
 label-switching
 enable-ldp ipv4

router ospf 1
```

```
 bfd all-interfaces
 network 4.4.4.4/32 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0

evpn mpls enable
evpn mpls vtep-ip-global 4.4.4.4
!
router bgp 100
neighbor 1.1.1.1 remote-as 100
neighbor 1.1.1.1 update-source lo
address-family l2vpn evpn
neighbor 1.1.1.1 activate
exit
!
mac vrf vrf2
rd 4.4.4.4:2
route-target both 2:2
!
evpn mpls id 2 xconnect target-mpls-id 52
host-reachability-protocol evpn-bgp vrf2
evi-name test2
!
interface xe1
switchport
load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
  access-if-evpn
   map vpn-id 52
ethernet cfm domain-type character-string domain-name 12346 level 7
mipcreation none
service ma-type string ma-name 124
  ethernet cfm mep up mpid 10 active true xe1.2001 vlan 2028
    cc multicast state enable
    ethernet cfm loss-measurement reply lmm
  ethernet cfm delay-measurement reply dmm
  exit-ether-ma-mep-mode
rmep auto-discovery enable
cc interval 10ms
exit-ether-ma-mode
```

P1:

```
interface lo
 ip address 2.2.2.2/32 secondary

router ldp
 transport-address ipv4 2.2.2.2

interface xe2
 ip address 192.168.10.2/24
 label-switching
 enable-ldp ipv4

interface xe1
 ip address 192.168.20.1/24
 label-switching
```

```
 enable-ldp ipv4

router ospf 1
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 192.168.10.0/24 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
```

P2:

```
interface lo
 ip address 3.3.3.3/32 secondary

router ldp
 transport-address ipv4 3.3.3.3

interface xe1
 ip address 192.168.20.2/24
 label-switching
 enable-ldp ipv4

interface xe2
 ip address 192.168.30.1/24
 label-switching
 enable-ldp ipv4

router ospf 1
 ospf router-id 3.3.3.3
 bfd all-interfaces
 network 3.3.3.3/32 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0
```

Validation

Verify the EVPN xconnect status.

```
PE1#show evpn mpls xconnect
EVPN Xconnect Info
======================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up
Local               Remote                        Connection-Details
================================================================================
VPN-ID      EVI-Name    MTU VPN-ID    Source Destination
PE-IP        MTU      Type       NW-Status
================================================================================
2 test2 1500 52 xe1.2001 --- Single Homed Port ---
4.4.4.4 1500 AC-NW NW-SET
```

Verify the CFM Errors:

```
PE1#show ethernet cfm errors domain 12346

Domain Name    MA Name    Level    VLAN       MEPID       Defects
-----------------------------------------------------------------
12346          124         7       2028        20          .....
```

Verify the RMEP is learned or not.

```
PE1#show ethernet cfm maintenance-points remote domain 12346
MA_NAME    MEPID    RMEPID    LEVEL    Rx CCM    RDI     PEER-MAC     TYPE
-------------------------------------------------------------------------------
124        20        10          7      Yes     False e8c5.7ae3.37ee  Learnt
```

Verify the Ping:

```
PE1#ping ethernet mac   e8c5.7ae3.37ee  unicast source 20 domain 12346 ma 124
 success rate is  100 (5/5)
```

Verify the local whether Local MEP is installed or not:

```
PE1#show ethernet cfm maintenance-points local mep domain 12346 ma-name 126
MPID Dir Lvl VLAN CC-Stat HW-Status CC-Intvl MAC-Address Def Port MD Name
-------------------------------------------------------------------------------
124 Up 7 2028 Enable Installed 10 ms e8c5.7afe.fae9 F xe1.2001 12346
```

Verify the ethernet cfm ma status domain is active or not.

```
PE1#show ethernet cfm ma status domain 12346 ma-name 124
MA NAME      STATUS
124          Active
```

Verify the Ping:

```
PE1#ping ethernet mac e8c5.7ae3.37ee unicast source 20 domain 12346 ma 124
success rate is 100 (5/5)
```

Verify the Traceroute:

```
PE1#traceroute ethernet e8c5.7ae3.37ee mepid 20 domain 12346 ma 124
MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
e8c5.7ae3.37ee 1 RlyHit Ingress IngOK
```

Verify the Delay-measurement:

```
PE1#delay-measurement type proactive profile-name DM rmep 10 mep 20 domain 12346 ma 124
PE1#2024 Apr 10 13:35:37.236 : PE1: ONMD : INFO : [CFM_PM_SESSION_INFO_5]: CFM Frame
Delay Measurement session started for MEP Id 20 and RMEP Id 10
PE2-7033#show ethernet cfm delay-measurement mep 20 domain 12346 ma-name 124
MD : 12346
MA : 124
MEP : 20
VLAN ID : 10
```

```
Interface : po1000.10
Peer MAC Address : 00cc.dd00.0000
CURRENT:
=====================
RMEP ID : 10
Measurement ID : 1
Measurement Type : DMM
Elapsed time(sec) : 53
Start Time : 2024 Apr 10 13:35:37
Suspect Flag : FALSE
Min Frame Delay(usec) : 19
Max Frame Delay(usec) : 20
Avg Frame Delay(usec) : 19
Min Inter FD Variation(usec): 0
Max Inter FD Variation(usec): 1
Avg Inter FD Variation(usec): 0
FRAME DELAY BINS
Bin Number       Bin Threshold(usec)        Bin Counter
=========================================================
1               0 - < 4999                    52
2               5000 - < 9999                 0
3               10000 - < 4294967295          0
INTER-FRAME DELAY BINS
Bin Number       Bin Threshold(usec)       Bin Counter
=========================================================
1                0 - < 4999                   51
2                5000 - < 4294967295          0
```

Verify the Loss-measurement:

```
PE1#loss-measurement type proactive profile-name LM rmep 10 mep 20 domain 12346 ma 124
PE1#2024 Apr 10 13:35:05.345 : PE1 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame Loss
Measurement started for MEP:20 MA:124 MD:12346
PE1#show ethernet cfm loss-measurement mep 20 domain 12346 ma-name 124

MEP: 20 MA: 124
CURRENT:
 Measurement ID : 1
  Suspect                    : False
  Measurement Type           : lmm
  Elapsed time(sec)          : 55
  Start Time                 : 2024 Apr 10 13:37:05
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
Far End frame loss ratio : 0
HISTORY:
Measurement ID : 1
Suspect : FALSE
```

```
Measurement Type : lmm
Elapsed time(sec) : 60
End Time : 2024 Apr 10 13:36:05
Near End loss : 0
Far End loss : 0
Near End accumulated loss : 0
Far End accumulated loss : 0
Near End frame loss ratio : 0
Far End frame loss ratio : 0
Near End frame loss ratio min : 0
Far End frame loss ratio min : 0
Near End frame loss ratio max : 0
Far End frame loss ratio max : 0
```

Verify the Synthetic Loss Measurement:

```
PE1#loss-measurement type proactive profile-name SLM rmep 10 mep 20 domain 12346 ma 124
PE1#2024 Apr 10 13:40:15.587 : PE1 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame Loss
Measurement  started for MEP:20 MA:124 MD:12346
PE1#show ethernet cfm loss-measurement mep 20 domain 12346 ma-name 124
MEP: 20 MA: 124
CURRENT:
 Measurement ID : 2
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 17
  Start Time                 : 2024 Apr 10 13:41:15
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
HISTORY:
 Measurement ID : 1
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 60
  End Time                   : 2024 Apr 10 13:41:15
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
  Near End frame loss ratio min  : 0
  Far End frame loss ratio min  : 0
  Near End frame loss ratio max  : 0
  Far End frame loss ratio max  : 0
```

Implementation Examples

Enterprise Connectivity Monitoring:

Scenario: A large enterprise operates multiple branch offices connected via Ethernet services provided by a service provider network.

Use Case: Y.1731 CFM over sub-interface using Single Home EVPN ELINE enables the enterprise to monitor theconnectivity and performance of its branch office connections. It facilitates proactive fault detection and management, ensuring reliable and uninterrupted communication between the headquarters and branch offices.

Service Provider Network Operations:

Scenario: A service provider manages a diverse range of Ethernet services for its enterprise customers, including VPNs, Internet access, and cloud connectivity.

Use Case: Y.1731 CFM over sub-interface using Single Home EVPN ELINE empowers the service provider to deliverhigh-quality Ethernet services with enhanced fault management capabilities. It enables the provider to quickly identify and resolve connectivity issues, minimize service downtime, and maintain customer satisfaction.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Y.1731	A standard defined by the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) that specifies performance monitoring and fault management for Ethernet-based networks.
Sub-interface	A logical division of a physical interface, typically used to separate traffic based on VLANs or other criteria. In this context, sub-interfaces are employed to establish distinct connections within the EVPN ELINE SH topology.
EVPN	Ethernet Virtual Private Network (VPN) is a technology that enables the creation of virtual private networks over an Ethernet-based infrastructure. It provides multi-tenancy and allows for the segmentation of traffic in service provider networks.
ELINE	ELINE is a type of EVPN service that provides point-to-point Ethernet connectivity between two sites.
Single Home (SH)	Refers to the configuration where a Customer Edge device (CE) is connected to only one Provider Edge device (PE) within an EVPN setup. It contrasts with the multi-homed configuration, where a CE may be connected to multiple PEs.
Maintenance End Point (MEP)	MEP is a CFM entity that resides at the edge of a CFM domain. It is responsible for generating and transmitting CFM protocol packets to detect faults and collect performance data.
Maintenance Domain (MD)	MD is a logical grouping of MEPs within a CFM network. MEPs within the same MD can communicate with each other to perform CFM functions such as fault detection and performance monitoring.
Maintenance Association(MA)	MA is a collection of MEPs associated with a specific service or set of services. It defines the scope of CFM operations within a maintenance domain.

Maintenance Point Identifier (MPID)	MPID is a unique identifier assigned to each MEP within a maintenance association. It is used to distinguish between different MEPs within the same MA.
Service Level Measurement (SLM)	SLM is a CFM function used to measure the loss characteristics of a network path. It collects data on packet loss, delay, and jitter to assess the quality of service provided by the network.
Loopback Message Generation (LMM)	LMM is a CFM function used to test end-to-end connectivity by generating loopback messages. These messages are transmitted from a MEP and looped back to the same MEP to verify bidirectional communication.
Delay Measurement Message (DMM)	DMM is a CFM function used to measure the one-way delay of packets transmitted across a network. It helps assess the performance of the network in terms of packet delivery time.
Continuity Check (CC)	CC is a CFM function used to verify the continuity of a service or network path by periodically sending continuity check messages between MEPs. It helps detect connectivity faults such as link failures or misconfigurations.

CHAPTER 7 Y.1731 and CFM Over EVPN-ELINE Multi-home

Overview

The Multi Home EVPN ELINE Y.1731 CFM over Sub-interface feature enables the monitoring and management of Ethernet Virtual Private Network (EVPN) E-Line services using the Y.1731 Connectivity Fault Management (CFM) protocol over sub-interfaces. This feature enhances fault detection and performance monitoring capabilities for EVPN E-Line services, allowing network operators to ensure high availability and reliability of their networks. By extending Y.1731 CFM functionality to sub-interfaces in multi home EVPN E-Line deployments, this feature provides comprehensive end-to-end visibility and control, enabling proactive fault detection, isolation, and troubleshooting.

CFM multi-homing allows Customer Edge (CE) device to connect more than one Provider Edge (PE) device. Multi-homing ensures redundant connectivity. The redundant PE device ensures that there is no traffic disruption when there is a network failure.

Feature Characteristics

- Utilizes sub-interfaces to partition Ethernet traffic within the Multi home EVPN ELINE architecture, enabling efficient service delivery and management.
- Implements EVPN ELINE architecture with multi-homing capabilities, facilitating the creation of Ethernet Virtual Private Networks with simplified configurations and reduced complexity.
- Provides robust fault detection mechanisms to identify connectivity issues, link failures, and service disruptions in Ethernet networks.

Benefits

- Provides detailed insights into Ethernet service performance, enabling proactive monitoring and optimization of network resources.
- Minimizes service downtime by promptly detecting and resolving faults, ensuring uninterrupted service delivery and customer satisfaction.
- Optimizes network resource utilization and bandwidth allocation by identifying and addressing connectivity issues in a timely manner.
- Facilitates rapid fault identification and isolation, accelerating troubleshooting processes and reducing mean time to repair (MTTR).

Ensures compliance with Service Level Agreements (SLAs) by maintaining service quality metrics within defined thresholds and objectives.

Configuration

Configure Multi Home EVPN ELINE Y.1731 CFM over Sub-interface for enhanced fault management in EVPN networks.

Topology

The following topology consists of customer edge routers CE1 and CE2 with IPv2 Provider Edge routers PE1 and PE2. These are interconnected through the core router P in the IPv4 MPLS provider networks.

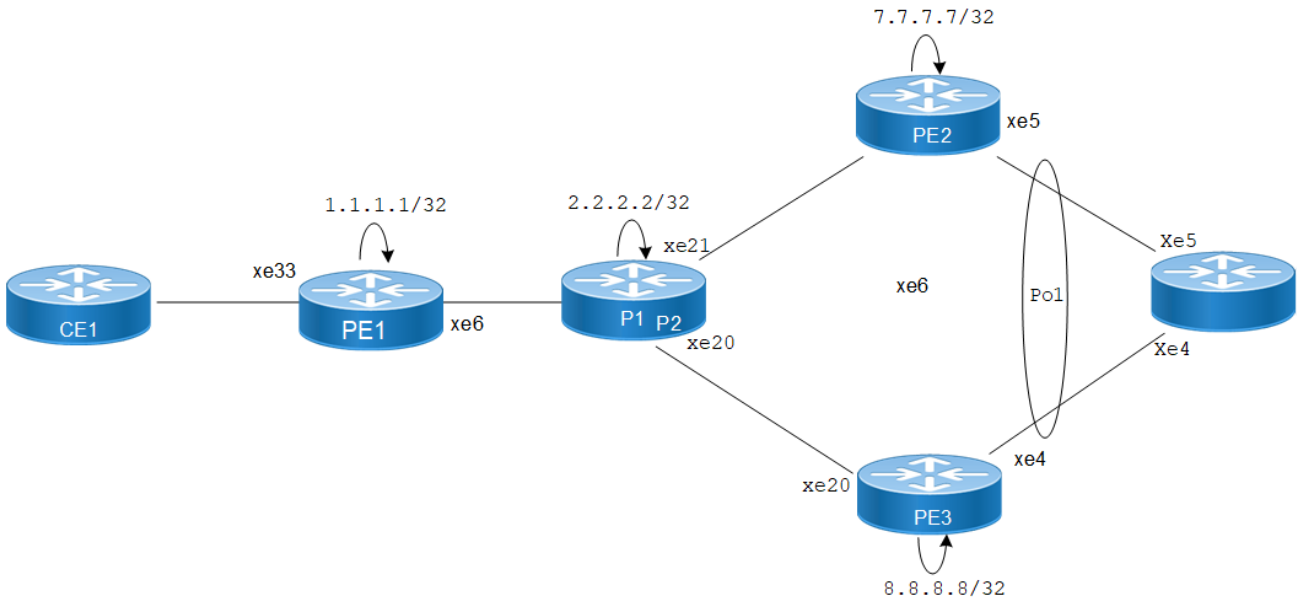

Figure 7-5: EVPN ELINE Over CFM Sub-interface

The following sessions displays the detailed information about configurations, and validations for CFM over sub-interface.

1. Configure Loopback Interface on PE1.

```
PE1(#configure terminal
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32
PE1(config-if)#exit
PE1(config-if)#commit
```

2. Configure Global LDP for distributing MPLS labels in the network.

```
PE1(config)# router ldp
PE1(config-router)# router-id 1.1.1.1
PE1(config-router)# targeted-peer ipv4 7.7.7.7
PE1(config-router)# targeted-peer ipv4 8.8.8.8
PE1(config-router-targeted-peer)#exit
PE1(config-router)# exit
PE1(config)# commit
```

3. Enable EVPN over MPLS and set a global VTEP IP.

```
PE1(config)# evpn mpls enable
PE1(config)# commit
PE1(config)# evpn mpls vtep-ip-global 1.1.1.1
PE1(config)# commit
```

4. Configure the interfaces connecting to the network, enabling LDP and MPLS label switching.

```
PE1(config)# interface xe6
PE1(config-if)# ip address 10.1.0.1/16
PE1(config-if)# enable-ldp ipv4
PE1(config-if)# label-switching
```

```
PE1(config-if)# exit
PE1(config)# commit
```

5. Set up OSPF for IP routing within the network.

```
PE1(config)# router ospf 1
PE1(config-router)# ospf router-id 1.1.1.1
PE1(config-router)# network 1.1.1.1/32 area 0
PE1(config-router)# network 10.1.0.0/16 area 0
PE1(config-router)# exit
PE1(config)# commit
```

6. Set up BGP for EVPN to exchange MAC and IP information.

```
PE1(config)# router bgp 1
PE1(config-router)# neighbor 7.7.7.7 remote-as 1
PE1(config-router)# neighbor 7.7.7.7 update-source lo
PE1(config-router)# neighbor 8.8.8.8 remote-as 1
PE1(config-router)# neighbor 8.8.8.8 update-source lo
PE1(config-router)# address-family l2vpn evpn
PE1(config-router-af)# neighbor 7.7.7.7 activate
PE1(config-router-af)# neighbor 8.8.8.8 activate
PE1(config-router-af)# exit
PE1(config-router)# exit
PE1(config)# commit
```

7. Configure MAC VRF.

```
PE1(config)# mac vrf vrf2
PE1(config-vrf)# rd 1.1.1.1:2
PE1(config-vrf)# route-target both 2:2
PE1(config-vrf)# exit
PE1(config)# commit
```

8. Configure EVPN and map VRF.

```
PE1(config)# evpn mpls id 52 xconnect target-mpls-id 2
PE1(config-evpn-mpls)# host-reachability-protocol evpn-bgp vrf2
PE1(config)# commit
```

9. Configure access port on interface xe33.2

```
PE1(config-if)# interface xe33.2 switchport
PE1(config-if)# description access-side-int
PE1(config-if)# encapsulation dot1q 2
PE1(config-if)# access-if-evpn
PE1(config-access-if)# map vpn-id 52
PE1(config-access-if)# exit
PE1(config)# commit
```

10. Set up CFM to monitor connectivity within the network.

```
PE1(config)#hardware-profile filter cfm-domain- name-str enable
PE1(config)# ethernet cfm domain-type character-string domain-name 12346 level 7
mip-creation none
PE1(config-ether-cfm-mpls-md)# service ma-type string ma-name 124
PE1(config-ether-cfm-mpls-ma)# ethernet cfm mep up mpid 10 active true xe33.2
vlan 2
PE1(config-ether-cfm-mpls-ma-mep)#cc multicast state enable
PE1(config-ether-cfm-mpls-ma-mep)#exit-ether- ma-mep-mode
PE1(config-ether-cfm-mpls-ma)#  rmep auto-discovery enable
PE1(config-ether-cfm-mpls-ma)#cc interval 10ms
PE1(config-ether-cfm-mpls-ma)#exit-ether-ma- mode
PE1(config-ether-cfm-mpls)#exit
```

```
PE1(config)#exit
PE1(config)#commit
```

Note: Similarly follow the same steps to configure respective `cfm mep up` and other CFM features for PE2 and PE3.

Configuration Snapshot:

PE1:

```
!
interface lo
 ip address 1.1.1.1/32
!
router ldp
 router-id 1.1.1.1
 targeted-peer 7.7.7.7
 targeted-peer 8.8.8.8
!
router ospf 1
 router-id 1.1.1.1
 network  1.1.1.1/32 area 0
 network  10.1.0.0/16 area 0
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 7.7.7.7 remote-as 1
 neighbor 7.7.7.7 update-source lo
 neighbor 8.8.8.8 remote-as 1
 neighbor 8.8.8.8 update-source lo
 !
 address-family l2vpn evpn
 neighbor 7.7.7.7 activate
 neighbor 8.8.8.8 activate
 exit-address-family
!
evpn mpls enable
evpn mpls vtep-ip-global 1.1.1.1
hardware-profile filter cfm-domain-name-str enable
hardware-profile statistics cfm-ccm enable
!
interface xe6
ip address 10.1.0.1/16
enable-ldp ipv4
label-switching
!
vrf definition vrf2
 rd 1.1.1.1:2
 route-target both 2:2
!
evpn mpls id 52 xconnect target-mpls-id 2
 host-reachability-protocol evpn-bgp vrf2
!
interface xe33.2 swicthport
 description access-side-int
 encapsulation dot1q 2
 access-if-evpn
```

```
  map vpn-id 52
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
  ethernet cfm mep up mpid 10 active true xe33.2 vlan 2
   cc multicast state enable
   exit-ether-ma- mode
   rmep auto-discovery enable
   cc interval 10ms
   exit-ether-ma- mode
!
```

P:

```
!
interface lo
 ip address 2.2.2.2/32
!
interface xe6
ip address 10.1.0.2/16
enable-ldp ipv4
label-switching
!
interface xe21
 ip address 123.1.1.1/24
 enable-ldp ipv4
label-switching
!
interface xe20
 ip address 124.1.1.1/24
 enable-ldp ipv4
label-switching
!
router ldp
 router-id 2.2.2.2
!
router ospf 1
 router-id 2.2.2.2
 network 2.2.2.2/32 area 0
 network 10.1.0.0/16 area 0
 network 123.1.1.0/24 area 0
 network 124.1.1.0/24 area 0
!
```

PE2:

```
!
interface lo
 ip address 7.7.7.7/32
!
interface xe21
 ip address 123.1.1.2/24
 enable-ldp ipv4
label-switching
!
router ldp
 router-id 7.7.7.7/32
 targeted-peer ipv4 1.1.1.1
```

```
 targeted-peer ipv4 8.8.8.8
!
router ospf 1
 router-id 7.7.7.7
 network  7.7.7.7/32 area 0
 network 123.1.1.0/24 area 0
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 update-source lo
 neighbor 8.8.8.8 remote-as 1
 neighbor 8.8.8.8 update-source lo
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 8.8.8.8 activate
 exit-address-family
!
evpn mpls enable
evpn mpls vtep-ip-global 7.7.7.7
hardware-profile filter evpn-mpls-mh enable
evpn mpls multihoming enable
!
vrf definition vrf2
 rd 7.7.7.7:2
 route-target both 2:2
!
interface Po1
 load-interval 30
 evpn multi-homed system-mac 0000.aaaa.bbbc
!
interface Po1.2 switchport
 encapsulation dot1q 2
 access-if-evpn
  map vpn-id 2
!
interface xe5
 channel-group 1 mode active
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
  ethernet cfm mep up mpid 20 active true po1.2 vlan 2
   cc multicast state enable
    ethernet cfm loss-measurement reply slm
    ethernet cfm delay-measurement reply dmm
   exit-ether-ma- mode
   rmep auto-discovery enable
   cc interval 10ms
   exit-ether-ma- mode
!
```

PE3:

```
!
interface lo
 ip address 8.8.8.8/32
!
```

```
interface xe20
 ip address 124.1.1.2/24
 enable-ldp ipv4
label-switching
!
interface xe4
 channel-group 1 mode active
!
router ldp
 router-id 8.8.8.8
 targeted-peer ipv4 1.1.1.1
 targeted-peer ipv4 7.7.7.7
!
router ospf 1
 router-id 8.8.8.8
 network  8.8.8.8/32 area 0
 network network 124.1.1.0/24 area 0
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 update-source lo
 neighbor 7.7.7.7 remote-as 1
 neighbor 7.7.7.7 update-source lo
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 7.7.7.7 activate
 exit-address-family
!
evpn mpls enable
evpn mpls vtep-ip-global 8.8.8.8
hardware-profile filter evpn-mpls-mh enable
evpn mpls multihoming enable
!
vrf definition vrf2
 rd 8.8.8.8:2
 route-target both 2:2
!
interface Po1
 load-interval 30
 evpn multi-homed system-mac 0000.aaaa.bbbc
!
interface Po1.2 switchport
 encapsulation dot1q 2
 access-if-evpn
  map vpn-id 2
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
  ethernet cfm mep up mpid 30 active true po1.2 vlan 2
   cc multicast state enable
    ethernet cfm loss-measurement reply slm
    ethernet cfm delay-measurement reply dmm
   exit-ether-ma- mode
   rmep auto-discovery enable
   cc interval 10ms
```

```
        exit-ether-ma- mode
        !
```

Validation

The following are the validations for PE1 and PE2.

PE1

The following validation is for PE1.

```
PE1#SH evpn mpls xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up



Local                               Remote       Connection-Details
=============================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
=============================================================================================
52           ----          1500  2            xe33.2       00:00:00:aa:aa:bb:bb:00:00:00
7.7.7.7         1500  AC-NW  NW-SET

8.8.8.8         1500  ----   ----
PE1#show ethernet cfm errors domain 12346

Domain Name         Level     MEPID     Defects
---------------------------------------------
12346               7         20        .....

PE1#show ethernet cfm ma status domain 12346 ma-name 124
MA NAME                  STATUS
------------------------------
    124                  Active
PE1#show ethernet cfm maintenance-points remote domain 12346 ma-name 124
 MEPID     RMEPID     LEVEL     Rx CCM    RDI    PEER-MAC        TYPE
-------------------------------------------------------------------------
 10        20         7         Yes       False  00aa.bb00.0002 Learnt
 10        30         7         Yes       False  00aa.dd00.0003 Learnt


PE1#show ethernet cfm maintenance-points local mep domain 12346 ma-name 124
MPID Dir Lvl CC-Stat HW-Status  CC-Intvl MAC-Address    Def Port  MD Name
```

```
--------------------------------------------------------------------------------
10   Up  7   Enable  Installed  100 ms   3417.ebe4.af22 F   xe33.2 12346

PE1#ping ethernet mac 00aa.bb00.0002 unicast source 10 domain 12346 ma 124
 success rate is  100 (5/5)


PE1#traceroute ethernet 00aa.bb00.0002 mepid 10 domain 12346 ma 124
MP Mac         Hops  Relay-action          Ingress/Egress  Ingress/Egress action
00aa.bb00.0002   1     RlyHit                 Ingress         IngOK

PE1#ping ethernet mac 00aa.dd00.0003 unicast source 10 domain 12346 ma 124
 success rate is  100 (5/5)
PE1-7011#traceroute ethernet 00aa.dd00.0003 mepid 10 domain 12346 ma 124
MP Mac         Hops  Relay-action          Ingress/Egress  Ingress/Egress action
00aa.dd00.0003   1     RlyHit                 Ingress         IngOK
```

Verify Delay Measurement:

```
PE1#delay-measurement type proactive profile-name DM rmep 20 mep 10 domain 12346 ma 124
PE1-7011#2019 Feb 14 10:34:53.935 : PE2-7033 : ONMD : INFO : [CFM_PM_SESSION_INFO_5]:
CFM Frame Delay Measurement session started for MEP Id 10 and RMEP Id 20

PE1#show ethernet cfm delay-measurement mep 10 domain 12346 ma-name 124
 MD                     : 12346
 MA                     : 124
 MEP                    : 10
 VLAN ID                : 2
 Interface              : xe33.2
 Peer MAC Address       : 00aa.bb00.0002

CURRENT:
======================
 RMEP ID        : 20
 Measurement ID : 3
 Measurement Type            : DMM
 Elapsed time(sec)           : 16
 Start Time                  : 2019 Feb 14 10:36:53
 Suspect Flag                : FALSE
 Min Frame Delay(usec)       : 23
 Max Frame Delay(usec)       : 24
 Avg Frame Delay(usec)       : 23
 Min Inter FD Variation(usec): 0
 Max Inter FD Variation(usec): 1
 Avg Inter FD Variation(usec): 0

             FRAME DELAY BINS
 Bin Number    Bin Threshold(usec)          Bin Counter
=========================================================
 1             0          - < 4999           16
 2             5000       - < 9999           0
 3             10000      - < 4294967295     0
```

```
           INTER-FRAME DELAY BINS
 Bin Number    Bin Threshold(usec)          Bin Counter
=========================================================
 1             0          - < 4999           15
 2             5000       - < 4294967295     0


HISTORY STATISTICS
======================
 MD                     : 12346
 MA                     : 124
 MEP                    : 10
 VLAN ID                : 2
 Interface              : xe33.2
 RMEP ID        : 20
 Measurement ID              : 1
 Measurement Type            : DMM
 Elapsed time(sec)           : 60
 End Time                    : 2019 Feb 14 10:35:53
 Suspect Flag                : FALSE
 Min Frame Delay(usec)       : 23
 Max Frame Delay(usec)       : 24
 Avg Frame Delay(usec)       : 23
 Min Inter FD Variation(usec): 0
 Max Inter FD Variation(usec): 1
 Avg Inter FD Variation(usec): 0

             FRAME DELAY BINS
 Bin Number    Bin Threshold(usec)          Bin Counter
=========================================================
 1             0          - < 4999           59
 2             5000       - < 9999           0
 3             10000      - < 4294967295     0

           INTER-FRAME DELAY BINS
 Bin Number    Bin Threshold(usec)          Bin Counter
=========================================================
 1             0          - < 4999           58
 2             5000       - < 4294967295     0

 RMEP ID        : 20
 Measurement ID              : 2
 Measurement Type            : DMM
 Elapsed time(sec)           : 60
 End Time                    : 2019 Feb 14 10:36:53
 Suspect Flag                : FALSE
 Min Frame Delay(usec)       : 23
 Max Frame Delay(usec)       : 24
 Avg Frame Delay(usec)       : 23
```

```
 Min Inter FD Variation(usec): 0
 Max Inter FD Variation(usec): 1
 Avg Inter FD Variation(usec): 0

              FRAME DELAY BINS
 Bin Number    Bin Threshold(usec)          Bin Counter
=========================================================
 1             0          - < 4999           60
 2             5000       - < 9999           0
 3             10000      - < 4294967295     0

           INTER-FRAME DELAY BINS
 Bin Number    Bin Threshold(usec)          Bin Counter
=========================================================
 1             0          - < 4999           59
 2             5000       - < 4294967295     0
```

Verify Synthetic Loss Measurement:

```
PE1#loss-measurement type proactive profile-name SLM rmep 20 mep 10 domain 12346 ma 124
PE1#2019 Feb 14 10:35:17.758 : PE2-7011 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame
Loss Measurement  started for MEP:10 MA:124 MD:12346
PE1-7011#show ethernet cfm loss-measurement mep 10 domain 12346 ma-name 124
MEP: 10 MA: 124
CURRENT:
 Measurement ID : 3
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 19
  Start Time                 : 2019 Feb 14 10:37:16
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0

HISTORY:

 Measurement ID : 1
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 60
  End Time                   : 2019 Feb 14 10:36:16
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
  Near End frame loss ratio min  : 0
```

```
  Far End frame loss ratio min  : 0
  Near End frame loss ratio max  : 0
  Far End frame loss ratio max  : 0

 Measurement ID : 2
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 60
  End Time                   : 2019 Feb 14 10:37:16
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
  Near End frame loss ratio min  : 0
  Far End frame loss ratio min  : 0
  Near End frame loss ratio max  : 0
  Far End frame loss ratio max  : 0
```

PE2/PE3

The following validations for PE2 and PE3.

```
PE2#show evpn mpls xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                              Remote       Connection-Details
================================ ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================================================================
2            ----          1500  52           po1.2        --- Single Homed Port ---
1.1.1.1         1500  AC-NW  NW-SET
PE2#show ethernet cfm errors domain 12346

Domain Name         Level     MEPID     Defects
---------------------------------------------
12346               7         20        .....
PE2#show ethernet cfm ma status domain 12346 ma-name 124
MA NAME                  STATUS
------------------------------
    124                  Active
```

```
PE2#show ethernet cfm maintenance-points local mep domain 12346 ma-name 124
MPID Dir Lvl CC-Stat HW-Status  CC-Intvl MAC-Address    Def Port  MD Name
-------------------------------------------------------------------------------
20   Up  7   Enable  Installed  100 ms   00aa.bb00.0002 F   po1.2 12346
PE2#show ethernet cfm maintenance-points remote domain 12346 ma-name 124
 MEPID     RMEPID     LEVEL     Rx CCM    RDI    PEER-MAC        TYPE
-------------------------------------------------------------------------
 20        10         7         Yes       False  3417.ebe4.af22 Learnt

PE2#ping ethernet mac 3417.ebe4.af22 unicast source 10 domain 12346 ma 124
 success rate is  100 (5/5)
PE2#traceroute ethernet 3417.ebe4.af22 mepid 10 domain 12346 ma 124
 MP Mac         Hops  Relay-action          Ingress/Egress  Ingress/Egress action
3417.ebe4.af22   1     RlyHit                Ingress         IngOK
```

CHAPTER 8 Y.1731 and CFM Over VPWS Sub-interface

Overview

Y.1731 Connectivity Fault Management (CFM) over Layer 2 Virtual Private Wire Service (VPWS) is a protocol and technology combination used for fault management in Layer 2 VPN networks. It allows for the detection and management of faults, performance monitoring, and fault localization within a VPWS network.

Feature Characteristics

- Facilitates end-to-end fault management across the VPWS network, covering provider and customer edges.
- Supports multi-level fault management, allowing operators to define different levels of fault detection and management for different parts of the network.
- Y.1731 CFM includes performance monitoring capabilities, such as delay measurement and frame loss measurement, to monitor service quality parameters.
- The protocol supports loopback and link trace functions to identify and troubleshoot faults within the VPWS network.

Benefits

- Enables rapid detection and localization of faults within the VPWS network, minimizing downtime and service disruptions.
- Provides performance monitoring capabilities, allowing to track key performance indicators and ensure service quality.
- Enhances network visibility by providing detailed fault and performance monitoring data, aiding in network troubleshooting and maintenance.

Prerequisites

Ensure the network devices participating in the L2VPN VPWS setup support Y.1731 CFM functionality. This includes the Provider Edge (PE) and Customer Edge (CE) devices.

Configuration

Configure Y.1731 CFM over sub-interface using L2VPN VPWS by defining the CFM domain, configuring service MEPs and MAs, and setting up cross-connects between primary and backup interfaces.

Topology

The topology consists of two Customer Edge devices (CE1 and CE2) connected to two Provider Edge devices (PE1 and PE2) via sub-interfaces (xe11 and xe12). The Provider Edge devices are interconnected through Provider Devices (P1 and P2). Y.1731 ethernet CFM is configured over these sub-interfaces to monitor and manage ethernet connectivity between the CE devices, ensuring fault detection and performance monitoring across the service provider's network.

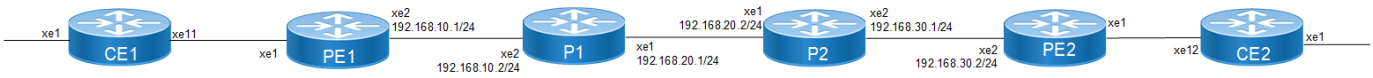

Figure 8-6: L2VPN VPWS Y1731 CFM Over Sub-interface

Perform the following configurations to configure Y.1731 CFM over sub-interface using L2VPN VPWS:

1. On Customer Edge (CE) Nodes (CE1 and CE2), configure the interface xe1 and set it as a switchport with a load interval of (`30 seconds`):

```
CE1(config)#interface xe1
CE1(config-if)#switchport
CE1(config-if)#load-interval 30
CE1(config-if)#commit
CE1(config-if)#exit
```

Note: Similarly follow the same steps to configure xe11(CE1) and xe12(CE2).

2. Create sub-interface (`xe1.2001`) adding the VLAN:

```
CE1(config)#interface xe1.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit

CE1(config)#interface xe11.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit
```

3. Set up a cross-connect named (`test100`), specifying in and out interfaces:

```
CE1(config)#cross-connect test100
CE1(config-xc)#interface xe1.2001
CE1(config-xc)#interface xe11.2001
CE1(config-xc)#commit
```

4. Perform the following on PE1:

    1. Configure CFM related hardware profiles:

```
PE1(config)# hardware-profile filter cfm-domain-name-str enable
PE1(config)# hardware-profile statistics cfm-lm enable
PE1(config)# hardware-profile statistics cfm-ccm enable
PE1(config)#hardware-profile statistics cfm-slm enable
```

    2. Configure the loopback interface with a secondary IP address(`1.1.1.1/32`):

```
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#commit
PE1(config-if)#exit
```

    3. Configure LDP targeted peers:

```
PE1(config)#router ldp
PE1(config-router)#targeted-peer ipv4 4.4.4.4
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#commit
PE1(config-router)exit
```

    4. Configure interface `xe2` with an IP address (`192.168.10.1/24`) and enable LDP:

```
PE1(config)#interface xe2
PE1(config-if)#load-interval 30
PE1(config-if)#ip address 192.168.10.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#commit
PE1(config-if)#exit
```

5. Configure OSPF routing, specify the OSPF router ID as (`1.1.1.1`), enable BFD on all interfaces, define the network (`1.1.1.1/32`) in area (`0.0.0.0`), and define the network (`192.168.10.0/24`) in area (`0.0.0.0`):

```
PE1(config)#router ospf 1
PE1(config-router)#ospf router-id 1.1.1.1
PE1(config-router)#bfd all-interfaces
PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0
PE1(config-router)#network 192.168.10.0/24 area 0.0.0.0
PE1(config-router)#commit
PE1(config-router)#exit
```

6. Set up an L2VPN pseudowire (`test1`) between PE1 and PE2.

```
PE1(config)#mpls l2-circuit test1 2001 4.4.4.4
PE1(config-pseudowire)#commit
PE1(config-pseudowire)#exit
```

7. Configure sub-interface (`xe1.2001`) as an access interface for VPWS.

```
PE1(config)#interface xe1.2001 switchport
PE1(config-if)#encapsulation dot1q 2028
PE1(config-if)#access-if-vpws
PE1(config-acc-if-vpws)#mpls-l2-circuit test1 primary
PE1(config-acc-if-vpws)#commit
PE1(config-acc-if-vpws)#exit
```

8. Configure Up-mep CFM domain:
   - Set the domain type as a character string with the domain name (`12346`) and (`level 7`)
   - Specify the MA type as a string with the MA name (`124`)
   - Associate the MA with (`VLAN 2028`)
   - Set up a MEP with MEP ID (`20`) as active on interface (`xe1.2001`)
   - Enable multicast state for continuity check, and auto-discovery of RMEPs
   - Set the continuity check interval to (`10 milliseconds`)

```
PE1(config)#ethernet cfm domain-type character-string domain-name
12346 level 7 mip-creation none
PE1(config-ether-cfm)# service ma-type string ma-name 124
PE1(config-ether-cfm-ma)#ethernet cfm mep up mpid 20 active true
xe1.2001 vlan 2028
PE1(config-ether-cfm-ma-mep)#cc multicast state enable
PE1(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
PE1(config-ether-cfm-ma)#rmep auto-discovery enable
PE1(config-ether-cfm-ma)#cc interval 10ms
PE1(config-ether-cfm-ma)#exit-ether-ma-mode
PE1(config-ether-cfm)#commit
PE1(config-ether-cfm)exit
```

   - Create a loss measurement profile named SLM with measurement type SLM, measurement interval of `1`, intervals stored of `3`, and message period of `(1)` second.

```
PE1(config)#ethernet cfm loss-measurement profile-name SLM
PE1(config-cfm-lm)#measurement-type slm
PE1(config-cfm-lm)#measurement-interval 1
PE1(config-cfm-lm)#intervals-stored 3
PE1(config-cfm-lm)#message-period 1s
PE1(config-cfm-lm)#exit
```

- Create loss measurement profile named LM with measurement type LMM, measurement interval of (`1`), intervals stored of (`3`), and message period of (`1 second`),

```
PE1(config)#ethernet cfm loss-measurement profile-name LM
PE1(config-cfm-lm)#measurement-type lmm
PE1(config-cfm-lm)#measurement-interval 1
PE1(config-cfm-lm)#intervals-stored 3
PE1(config-cfm-lm)#message-period 1s
PE1(config-cfm-lm)#exit
```

- Create a delay measurement profile named DM with a measurement interval of (`1`), intervals stored of (`2`), and message period of `(1 second)`.

```
PE1(config)#ethernet cfm delay-measurement profile-name DM
PE1(config-cfm-dm)#measurement-interval 1
PE1(config-cfm-dm)#intervals-stored 2
PE1(config-cfm-dm)#message-period 1
```

Configuration Snapshot:

CE1:

```
interface xe1
switchport
 load-interval 30

interface xe1.2001 switchport
 encapsulation dot1q 2028

interface xe11.2001 switchport
 encapsulation dot1q 2028

cross-connect test100
 interface xe1.2001
 interface xe11.2001
```

CE2:

```
interface xe1
switchport
 load-interval 30

interface xe1.2001 switchport
 encapsulation dot1q 2028

interface xe12.2001 switchport
 encapsulation dot1q 2028

cross-connect test100
 interface xe1.2001
 interface xe12.2001
```

PE1:

```
interface lo
```

```
 ip address 1.1.1.1/32 secondary

router ldp
 targeted-peer ipv4 4.4.4.4

interface xe2
 load-interval 30
 ip address 192.168.10.1/24
 label-switching
 enable-ldp ipv4

router ospf 1
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 192.168.10.0/24 area 0.0.0.0

mpls l2-circuit test1 2001 4.4.4.4

interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-vpws
 mpls-l2-circuit test1 primary

ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
   ethernet cfm mep up mpid 20 active true xe1.2001 vlan 2028
     cc multicast state enable
     exit-ether-ma-mep-mode
   rmep auto-discovery enable
   cc interval 10ms
   exit-ether-ma-mode

ethernet cfm loss-measurement profile-name SLM
 measurement-type slm
 measurement-interval 1
 intervals-stored 3
 message-period 1s
!
ethernet cfm loss-measurement profile-name LM
 measurement-type lmm
 measurement-interval 1
 intervals-stored 3
 message-period 1s
!
ethernet cfm delay-measurement profile-name DM
 measurement-interval 1
 intervals-stored 2
 message-period 1s
```

PE2:

```
interface lo
 ip address 4.4.4.4/32 secondary

router ldp
 targeted-peer ipv4 1.1.1.1
```

```
interface xe2
 load-interval 30
 ip address 192.168.30.2/24
 label-switching
 enable-ldp ipv4

router ospf 1
ospf router-id 4.4.4.4
 bfd all-interfaces
 network 4.4.4.4/32 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0

mpls l2-circuit test1 2001 1.1.1.1

interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-vpws
 mpls-l2-circuit test1 primary

ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
   ethernet cfm mep up mpid 10 active true xe1.2001 vlan 2028
     cc multicast state enable
     ethernet cfm loss-measurement reply lmm
     ethernet cfm delay-measurement reply dmm
     exit-ether-ma-mep-mode
   rmep auto-discovery enable
   cc interval 10ms
   exit-ether-ma-mode
```

P1:

```
interface lo
 ip address 2.2.2.2/32 secondary

router ldp
 transport-address ipv4 2.2.2.2

interface xe2
 ip address 192.168.10.2/24
 label-switching
 enable-ldp ipv4

interface xe1
 ip address 192.168.20.1/24
 label-switching
 enable-ldp ipv4

router ospf 1
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 192.168.10.0/24 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
```

P2:

```
interface lo
 ip address 3.3.3.3/32 secondary

router ldp
 transport-address ipv4 3.3.3.3

interface xe1
 ip address 192.168.20.2/24
 label-switching
 enable-ldp ipv4

interface xe2
 ip address 192.168.30.1/24
 label-switching
 enable-ldp ipv4

router ospf 1
 ospf router-id 3.3.3.3
 bfd all-interfaces
 network 3.3.3.3/32 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0
```

Validation

Verify the RMEP is learned or not.

```
PE1#show ethernet cfm maintenance-points remote domain 12346
 MA_NAME  MEPID     RMEPID     LEVEL     Rx CCM    RDI    PEER-MAC        TYPE
--------------------------------------------------------------------------------
 124       20        10         7         Yes       False  e8c5.7ae3.37ee Learnt
```

Verify the CFM Errors:

```
PE1#show ethernet cfm errors domain 12346

Domain Name         Level     MEPID     Defects
-----------------------------------------------
12346               7         20        .....


1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

Verify the CFM status:

```
PE1#show ethernet cfm ma status domain 12346 ma-name 124
MA NAME                  STATUS
124                  Active
```

Verify the Ping:

```
PE1#ping ethernet mac   e8c5.7ae3.37ee  unicast source 20 domain 12346 ma 124
 success rate is  100 (5/5)
```

Verify the Traceroute:

```
PE1#traceroute ethernet   e8c5.7ae3.37ee mepid 20 domain 12346 ma 124
 MP Mac         Hops  Relay-action          Ingress/Egress  Ingress/Egress action
e8c5.7ae3.37ee   1     RlyHit                 Ingress         IngOK
```

Verify the MPLS virtual circuit table, which contains information about MPLS label-switched paths (LSPs) and its associated virtual circuits in the network.

```
PE1#show mpls vc-table
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VC-ID      Vlan-ID  Inner-Vlan-ID  Access-Intf   Network-Intf  Out Label  Tunnel-Label
Nexthop         Status     UpTime
2001        N/A      N/A           xe1.2001        xe2        26240      25601
4.4.4.4         Active     00:38:02
```

Verify the Delay-measurement:

```
PE1#delay-measurement type proactive profile-name DM rmep 10 mep 20 domain 12346 ma 124
PE1#2023 Oct 12 04:11:56.696 : PE1 : ONMD : INFO : [CFM_PM_SESSION_INFO_5]: CFM Frame
Delay Measurement session started for MEP Id 20 and RMEP Id 10

PE1#show ethernet cfm delay-measurement mep 20 domain 12346 ma-name 124
 MD                    : 12346
 MA                    : 124
 MEP                   : 20
 VC Name               : test3
 Peer MAC Address      : e8c5.7ae3.37ee

CURRENT:

 RMEP ID        : 10
 Measurement ID : 1
 Measurement Type            : DMM
 Elapsed time(sec)           : 2
 Start Time                  : 2023 Oct 12 04:11:56
 Suspect Flag                : FALSE
 Min Frame Delay(usec)       : 40
 Max Frame Delay(usec)       : 74
 Avg Frame Delay(usec)       : 57
 Min Inter FD Variation(usec): 34
 Max Inter FD Variation(usec): 34
 Avg Inter FD Variation(usec): 34



         FRAME DELAY BINS
Bin Number    Bin Threshold(usec)          Bin Counter
 1             0          - < 4999          2
 2             5000       - < 9999          0
```

```
 3             10000      - < 14999          0
 4             15000      - < 4294967295     0



       INTER-FRAME DELAY BINS
Bin Number    Bin Threshold(usec)          Bin Counter
 1             0          - < 4999           1
 2             5000       - < 9999           0
 3             10000      - < 4294967295     0
```

Verify the Loss-measurement:

```
PE1#loss-measurement type proactive profile-name LM rmep 10 mep 20 domain 12346 ma 124
2023 Oct 12 04:18:43.667 : PE1 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame Loss
Measurement  started for MEP:20 MA:124 MD:12346
PE1#show ethernet cfm loss-measurement mep 20 domain 12346 ma-name 124

MEP: 20 MA: 124
CURRENT:
 Measurement ID : 1
  Suspect                    : False
  Measurement Type           : lmm
  Elapsed time(sec)          : 10
  Start Time                 : 2023 Oct 12 04:18:43
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
```

Verify the Synthetic Loss Measurement:

```
PE1#loss-measurement type proactive profile-name SLM rmep 10 mep 20 domain 12346 ma 124
PE1#2024 Apr 10 13:40:15.587 : PE1 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame Loss
Measurement  started for MEP:20 MA:124 MD:12346
PE1#show ethernet cfm loss-measurement mep 20 domain 12346 ma-name 124
MEP: 20 MA: 124
CURRENT:
 Measurement ID : 2
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 17
  Start Time                 : 2024 Apr 10 13:41:15
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
HISTORY:
```

```
Measurement ID : 1
 Suspect                    : False
 Measurement Type           : slm
 Elapsed time(sec)          : 60
 End Time                   : 2024 Apr 10 13:41:15
 Near End loss              : 0
 Far End loss               : 0
 Near End accumulated loss  : 0
 Far End accumulated loss   : 0
 Near End frame loss ratio  : 0
 Far End frame loss ratio   : 0
 Near End frame loss ratio min  : 0
 Far End frame loss ratio min  : 0
 Near End frame loss ratio max  : 0
 Far End frame loss ratio max  : 0
```

Verify the DM, LM, and SLM active sessions.

```
PE1#show ethernet cfm maintenance-points count

Total No of MIPs                : 0
Total No of MEPs                : 2
Total No of UP MEPs             : 2
Total No of Down MEPs           : 0
Total No of Active CCM sessions  : 2
Total No of UP CCM sessions      : 2
Total No of Active LM sessions   : 2
Total No of Active DM sessions   : 1
```

Implementation Examples

- To support a vast network infrastructure delivering VPWS to a multitude of enterprise clients, it is imperative to maintain uninterrupted connectivity and peak performance for these VPWS connections, all while minimizing the risk of downtime or disruptions.
- Understanding the role of fault detection, localization, and performance monitoring within the VPWS network, deploy Y.1731 CFM over Layer 2 VPN (VPWS) to enhance the network's resilience and operational efficiency.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Connectivity Fault Management (CFM)	CFM is a protocol used to detect, verify, and isolate connectivity faults in a network. It operates at the data link layer and is designed to monitor ethernet networks.

Virtual Private Wire Service (VPWS)	VPWS is a Layer 2 VPN service that provides point-to-point connectivity between two sites over an MPLS network. It emulates a leased line or circuit between the customer premises equipment (CPE) devices.
Maintenance End Point (MEP)	MEP is a CFM entity that resides at the edge of a CFM domain. It is responsible for generating and transmitting CFM protocol packets to detect faults and collect performance data.
Maintenance Domain (MD)	MD is a logical grouping of MEPs within a CFM network. MEPs within the same MD can communicate with each other to perform CFM functions such as fault detection and performance monitoring.
Maintenance Association(MA)	MA is a collection of MEPs associated with a specific service or set of services. It defines the scope of CFM operations within a maintenance domain.
Maintenance Point Identifier (MPID)	MPID is a unique identifier assigned to each MEP within a maintenance association. It is used to distinguish between different MEPs within the same MA.
Service Level Measurement (SLM)	SLM is a CFM function used to measure the loss characteristics of a network path. It collects data on packet loss, delay, and jitter to assess the quality of service provided by the network.
Loopback Message Generation (LMM)	LMM is a CFM function used to test end-to-end connectivity by generating loopback messages. These messages are transmitted from a MEP and looped back to the same MEP to verify bidirectional communication.
Delay Measurement Message (DMM)	DMM is a CFM function used to measure the one-way delay of packets transmitted across a network. It helps assess the performance of the network in terms of packet delivery time.
Continuity Check (CC)	CC is a CFM function used to verify the continuity of a service or network path by periodically sending continuity check messages between MEPs. It helps detect connectivity faults such as link failures or misconfigurations.

CHAPTER 9 Y.1731 and CFM Over EVPN ELAN Single Home

Overview

The Single Home EVPN ELAN Y.1731 CFM over Sub-interface feature enables the monitoring and management ofEthernet Virtual Private Network (EVPN) Ethernet LAN services using the Y.1731 Connectivity Fault Management (CFM) protocol over sub-interfaces. This feature enhances fault detection and performance monitoring capabilities for EVPN E-LAN services, allowing network operators to ensure high availability and reliability of their networks. By extending Y.1731 CFM functionality to sub-interfaces in single home EVPN E-LAN deployments, this feature provides comprehensive end-to-end visibility and control, enabling proactive fault detection, isolation, and troubleshooting.

Feature Characteristics

- Utilizes sub-interfaces to partition Ethernet traffic within the Single Home EVPN ELAN architecture, enabling efficient service delivery and management.
- Implements EVPN ELAN architecture with single-homing capabilities, facilitating the creation of Ethernet Virtual Private Networks with simplified configurations and reduced complexity.
- Provides robust fault detection mechanisms to identify connectivity issues, link failures, and service disruptions in Ethernet networks.

Benefits

- Provides detailed insights into Ethernet service performance, enabling proactive monitoring and optimization of network resources.
- Minimizes service downtime by promptly detecting and resolving faults, ensuring uninterrupted service delivery and customer satisfaction.
- Optimizes network resource utilization and bandwidth allocation by identifying and addressing connectivity issues in a timely manner.
- Facilitates rapid fault identification and isolation, accelerating troubleshooting processes and reducing mean time to repair (MTTR).
- Ensures compliance with Service Level Agreements (SLAs) by maintaining service quality metrics within defined thresholds and objectives.

Prerequisites

Ensure that the network devices (routers, switches) support Y.1731 CFM functionality and Single Home EVPN ELAN configuration.

Verify that the devices are running compatible software versions that include support for these features.

Configuration

Configure Single Home EVPN ELAN Y.1731 CFM over Sub-interface for enhanced fault management in EVPN

networks.

Topology

The topology consists of three Customer Edge devices (CE1, CE2, and CE3) connected to Provider Edge devices (PE1,PE2, and PE3) through sub-interfaces. The Provider Edge devices are interconnected through Provider devices (P1 and P2).
Y.1731 functionality is implemented over these sub-interfaces, allowing for fault detection and performance monitoring of Ethernet connectivity between the customer sites.

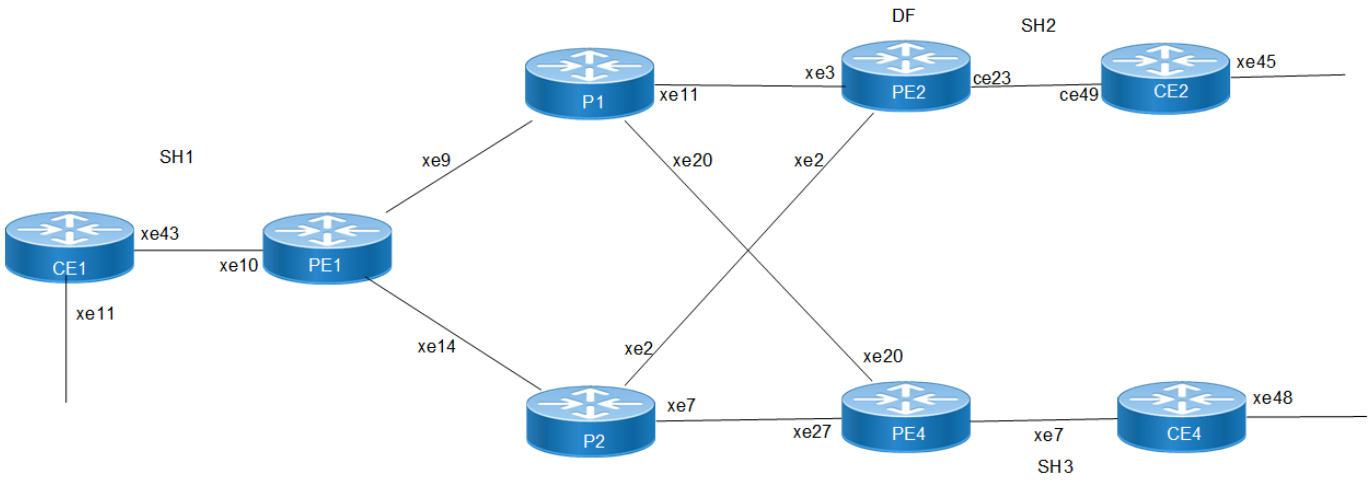

Figure 9-7: EVPN ELAN Over Sub-interface-Single Home

Perform the following configurations to configure Single Home EVPN ELAN Y.1731 CFM over Sub-interface:

1. On Customer Edge (CE) Nodes (CE1, CE2, and CE3), configure the interface xe1 and set it as a switchport with a load interval of (`30 seconds`):

```
CE1(config)#interface xe1
CE1(config-if)#switchport
CE1(config-if)#load-interval 30
CE1(config-if)#commit
CE1(config-if)#exit
```

Note: Similarly follow the same steps to configure xe11(CE1), xe12(CE2), and xe13(CE3).

2. Create sub-interface (`xe1.2001`) adding the VLAN:

```
CE1(config)#interface xe1.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit

CE1(config)#interface xe11.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit
```

3. Set up a cross-connect named (`test100`), specifying in and out interfaces:

```
CE1(config)#cross-connect test100
CE1(config-xc)#interface xe1.2001
CE1(config-xc)#interface xe11.2001
```

```
CE1(config-xc)#commit
```

4. Perform the following on PE1:

   1. Configure CFM related hardware profiles:

      ```
      PE1(config)# hardware-profile filter cfm-domain-name-str enable
      PE1(config)# hardware-profile statistics cfm-lm enable
      PE1(config)# hardware-profile statistics cfm-ccm enable
      PE1(config)#hardware-profile statistics cfm-slm enable
      ```

   2. Configure the loopback interface with a secondary IP address(`1.1.1.1/32`):

      ```
      PE1(config)#interface lo
      PE1(config-if)#ip address 1.1.1.1/32 secondary
      PE1(config-if)#commit
      PE1(config-if)#exit
      ```

   3. Configure LDP targeted peers:

      ```
      PE1(config)#router ldp
      PE1(config-router)#targeted-peer ipv4 4.4.4.4
      PE1(config-router-targeted-peer)#exit-targeted-peer-mode
      PE1(config-router)#commit
      PE1(config-router)exit
      ```

   4. Configure interface `xe2` with an IP address (`192.168.10.1/24`) and enable LDP:

      ```
      PE1(config)#interface xe2
      PE1(config-if)#load-interval 30
      PE1(config-if)#ip address 192.168.10.1/24
      PE1(config-if)#label-switching
      PE1(config-if)#enable-ldp ipv4
      PE1(config-if)#commit
      PE1(config-if)#exit
      ```

   5. Configure OSPF routing, specify the OSPF router ID as (`1.1.1.1`), enable BFD on all interfaces, define the network (`1.1.1.1/32`) in area (`0.0.0.0`), and define the network (`192.168.10.0/24`) in area (`0.0.0.0`):

      ```
      PE1(config)#router ospf 1
      PE1(config-router)#ospf router-id 1.1.1.1
      PE1(config-router)#bfd all-interfaces
      PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0
      PE1(config-router)#network 192.168.10.0/24 area 0.0.0.0
      PE1(config-router)#commit
      PE1(config-router)#exit
      ```

   6. Enable EVPN MPLS globally and configure VTEP IP:

      ```
      PE1(config)# evpn mpls enable
      PE1(config)# commit
      PE1(config)# evpn mpls vtep-ip-global 1.1.1.1
      PE1(config)# commit
      ```

   7. Configure BGP with the remote PE devices and activate EVPN:

      ```
      PE1(config)# router bgp 100
      PE1(config-router)# neighbor 4.4.4.4 remote-as 100
      PE1(config-router)# neighbor 4.4.4.4 update-source lo
      PE1(config-router)# address-family l2vpn evpn
      PE1(config-router-af)# neighbor 4.4.4.4 activate
      PE1(config-router-af)# exit
      PE1(config-router)# exit
      PE1(config)# commit
      ```

8. Configure MAC VRF with the appropriate RD and RT:

```
PE1(config)# mac vrf vrf2
PE1(config-vrf)# rd 1.1.1.1:2
PE1(config-vrf)# route-target both 2:2
PE1(config-vrf)# exit
```

9. Map the EVPN instance and VRF, specifying the EVPN ID:

```
PE1(config)#evpn mpls id 101
PE1(config-evpn-mpls)# host-reachability-protocol evpn-bgp vrf2
PE1(config-evpn-mpls)#commit
PE1(config-evpn-mpls)# commit
PE1(config-router-af)# exit
```

10. Configure access ports on PE1:

```
PE1(config)# interface xe1.2001 switchport
PE1(config-if)# encapsulation dot1q 2028
PE1(config-if)# access-if-evpn
PE1(config-acc-if-evpn)# map vpn-id 101
PE1(config-acc-if-evpn)# commit
```

11. Configure CFM MEP on PE1, define the FCM domain (`12346`), create MA, configure MEP, and configure Remote MEP Auto-discovery,set CC Interval `10ms`:

```
PE1(config)# ethernet cfm domain-type character-string domain-name12346
level 7 mip-creation default
PE1(config-ether-cfm)# service ma-type string ma-name 124
PE1(config-ether-cfm-ma)# ethernet cfm mep up mpid 20 active true
xe1.2001 vlan 2028
PE1(config-ether-cfm-ma-mep)# cc multicast state enable
PE1(config-ether-cfm-ma-mep)# exit-ether-ma-mep-mode
PE1(config-ether-cfm-ma)# rmep auto-discovery enable
PE1(config-ether-cfm-ma)# cc interval 10ms
PE1(config-ether-cfm-ma)# exit-ether-ma-mode
PE1(config-ether-cfm)# commit
```

12. Provide CFM configuration, define a delay measurement profile named DM, set the measurement interval to 1 second, specify the number of intervals stored as 2, configure the message period as 1 second, set the measurement type to LMM, set the measurement interval to 1 second, specify the number of intervals stored as 3, define a service level measurement profile named SLM, set the measurement type to SLM:

```
PE1(config)# ethernet cfm delay-measurement profile-name DM
PE1(config-cfm-dm)# measurement-interval 1
PE1(config-cfm-dm)# intervals-stored 2
PE1(config-cfm-dm)# message-period 1s
PE1(config-cfm-dm)# commit

PE1(config)# ethernet cfm loss-measurement profile-name SLM
PE1(config-cfm-lm)# measurement-type slm
PE1(config-cfm-lm)# measurement-interval 1
PE1(config-cfm-lm)# intervals-stored 3
PE1(config-cfm-lm)# message-period 1s
PE1(config-cfm-lm)# commit
```

Configuration Snapshot:

CE1:

```
interface xe1
switchport
load-interval 30
```

```
!
interface xe1.2001 switchport
encapsulation dot1q 2028
!

interface xe11.2001 switchport
encapsulation dot1q 2028
!
cross-connect test100
interface xe1.2001
interface xe11.2001
```

CE2:

```
interface xe1
switchport
load-interval 30
!
interface xe1.2001 switchport
encapsulation dot1q 2028
!
interface xe12.2001 switchport
encapsulation dot1q 2028
!
cross-connect test100
interface xe1.2001
interface xe12.2001
```

PE1:

```
interface lo
 ip address 1.1.1.1/32 secondary
!
router ldp
 targeted-peer ipv4 4.4.4.4
  exit-targeted-peer-mode
 targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
transport-address ipv4 1.1.1.1
!
interface xe2
 load-interval 30
 ip address 192.168.10.1/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 192.168.10.0/24 area 0.0.0.0
!
evpn mpls enable
evpn mpls vtep-ip-global 1.1.1.1
!
router bgp 100
 neighbor 4.4.4.4 remote-as 100
 neighbor 4.4.4.4 update-source lo
```

```
 neighbor 5.5.5.5 remote-as 100
 neighbor 5.5.5.5 update-source lo
 address-family l2vpn evpn
 neighbor 4.4.4.4 activate
 neighbor 5.5.5.5 activate
exit
!
mac vrf vrf2
 rd 1.1.1.1:2
 route-target both 2:2
!
evpn mpls id 101
 host-reachability-protocol evpn-bgp vrf2
!
interface xe1
 switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-evpn
 map vpn-id 101
!
ethernet cfm domain-type character-string domain-name 12346 level 7
mipcreation none
 service ma-type string ma-name 124
 ethernet cfm mep up mpid 20 active true xe1.2001 vlan 2028
 cc multicast state enable
 exit-ether-ma-mep-mode
 rmep auto-discovery enable
 cc interval 10ms
 exit-ether-ma-mode
!
ethernet cfm loss-measurement profile-name SLM
 measurement-type slm
 measurement-interval 1
 intervals-stored 3
 message-period 1s
!
ethernet cfm delay-measurement profile-name DM
 measurement-interval 1
 intervals-stored 2
 message-period 1s
```

PE2:

```
interface lo
 ip address 4.4.4.4/32 secondary
!
router ldp
 targeted-peer ipv4 1.1.1.1
exit-targeted-peer-mode
targeted-peer ipv4 5.5.5.5
exit-targeted-peer-mode
transport-address ipv4 4.4.4.4
!
interface xe2
 load-interval 30
```

```
 ip address 192.168.30.2/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id 4.4.4.4
 bfd all-interfaces
 network 4.4.4.4/32 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0
!
evpn mpls enable
evpn mpls vtep-ip-global 4.4.4.4
!
router bgp 100
 neighbor 1.1.1.1 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 5.5.5.5 remote-as 100
 neighbor 5.5.5.5 update-source lo
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 5.5.5.5 activate
exit
!
mac vrf vrf2
 rd 4.4.4.4:2
 route-target both 2:2
!
evpn mpls id 101
 host-reachability-protocol evpn-bgp vrf2
!
interface xe1
 switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-evpn
 map vpn-id 101
!
ethernet cfm domain-type character-string domain-name 12346 level 7
mipcreation none
 service ma-type string ma-name 124
 ethernet cfm mep up mpid 10 active true xe1.2001 vlan 2028
 cc multicast state enable
 ethernet cfm loss-measurement reply slm
 ethernet cfm delay-measurement reply dmm
 exit-ether-ma-mep-mode
 rmep auto-discovery enable
 cc interval 10ms
 exit-ether-ma-mode
!
```

PE3:

```
interface lo
 ip address 5.5.5.5/32 secondary
!
router ldp
```

```
 targeted-peer ipv4 1.1.1.1
exit-targeted-peer-mode
targeted-peer ipv4 4.4.4.4
exit-targeted-peer-mode
transport-address ipv4 5.5.5.5
!
interface xe3
 load-interval 30
 ip address 192.168.40.2/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id
 bfd all-interfaces
 network 5.5.5.5/32 area 0.0.0.0
 network 192.168.40.0/24 area 0.0.0.0
!
evpn mpls enable
evpn mpls vtep-ip-global 5.5.5.5
!
router bgp 100
 neighbor 1.1.1.1 remote-as 100
 neighbor 1.1.1.1 update-source lo
 neighbor 4.4.4.4 remote-as 100
 neighbor 4.4.4.4 update-source lo
 address-family l2vpn evpn
 neighbor 1.1.1.1 activate
 neighbor 4.4.4.4 activate
exit
!
mac vrf vrf2
 rd 5.5.5.5:2
 route-target both 2:2
!
evpn mpls id 101
 host-reachability-protocol evpn-bgp vrf2
!
interface xe1
 switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-evpn
  map vpn-id 101
!
ethernet cfm domain-type character-string domain-name 12346 level 7
mipcreation none
 service ma-type string ma-name 124
 ethernet cfm mep up mpid 30 active true xe1.2001 vlan 2028
 cc multicast state enable
  ethernet cfm loss-measurement reply slm
 ethernet cfm delay-measurement reply dmm
 exit-ether-ma-mep-mode
 rmep auto-discovery enable
 cc interval 10ms
```

```
 exit-ether-ma-mode
!
```

P1:

```
interface lo
 ip address 2.2.2.2/32 secondary

router ldp
 transport-address ipv4 2.2.2.2

interface xe2
 ip address 192.168.10.2/24
 label-switching
 enable-ldp ipv4

interface xe1
 ip address 192.168.20.1/24
 label-switching
 enable-ldp ipv4

router ospf 1
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 192.168.10.0/24 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
```

P2:

```
interface lo
 ip address 3.3.3.3/32 secondary

router ldp
 transport-address ipv4 3.3.3.3

interface xe1
 ip address 192.168.20.2/24
 label-switching
 enable-ldp ipv4

interface xe2
 ip address 192.168.30.1/24
 label-switching
 enable-ldp ipv4

router ospf 1
 ospf router-id 3.3.3.3
 bfd all-interfaces
 network 3.3.3.3/32 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0
```

CE3:

```
interface xe1
switchport
```

```
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
!
interface xe13.2001 switchport
 encapsulation dot1q 2028
!
cross-connect test100
 interface xe1.2001
 interface xe13.2001
```

Validation

Verify the EVPN MPLS status.

```
PE1#show evpn mpls
EVPN-MPLS Information
====================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VPN-ID   EVI-Name     EVI-Type Type Interface ESI                              VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
101      ----         L2       NW   ----      ----                             ----      -
---      1.1.1.1          4.4.4.4
101      ----         L2       NW   ----      ----                             ----      -
---      1.1.1.1          5.5.5.5
101      ----         --       AC   xe1.2001  --- Single Homed Port ---        ----      -
---      ----             ----
Total number of entries are 4
```

Verify the RMEP is learned or not:

```
PE1#show ethernet cfm maintenance-points remote domain 12346
 MA_NAME MEPID RMEPID LEVEL Rx CCM RDI PEER-MAC TYPE
-------------------------------------------------------------------------------
 124 20 10 7 Yes False e8c5.7ae3.37ee Learnt
 124 20 30 7 Yes False e8c5.7ae3.38ee Learnt
```

Verify the Ping:

```
PE1#ping ethernet mac e8c5.7ae3.37ee unicast source 20 domain 12346 ma 124
 success rate is 100 (5/5)

PE1#ping ethernet mac e8c5.7ae3.38ee unicast source 20 domain 12346 ma 124
 success rate is 100 (5/5))
```

Verify the Traceroute:

```
PE1#traceroute ethernet e8c5.7ae3.37ee mepid 20 domain 12346 ma 124
```

```
 MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
e8c5.7ae3.37ee 1 RlyHit Ingress IngOK

PE1#traceroute ethernet e8c5.7ae3.38ee mepid 20 domain 12346 ma 124
 MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
e8c5.7ae3.38ee 1 RlyHit Ingress IngOK
```

Verify the Delay-measurement:

```
PE1#delay-measurement type proactive profile-name DM rmep 10 mep 20 domain 12346 ma 124
PE1#2024 Apr 10 13:35:37.236 : PE1: ONMD : INFO : [CFM_PM_SESSION_INFO_5]: CFM Frame
Delay Measurement session started for MEP Id 20 and RMEP Id 10
PE2-7033#show ethernet cfm delay-measurement mep 20 domain 12346 ma-name 124
MD : 12346
MA : 124
MEP : 20
VLAN ID : 10
Interface : po1000.10
Peer MAC Address : 00cc.dd00.0000
CURRENT:
======================
RMEP ID : 10
Measurement ID : 1
Measurement Type : DMM
Elapsed time(sec) : 53
Start Time : 2024 Apr 10 13:35:37
Suspect Flag : FALSE
Min Frame Delay(usec) : 19
Max Frame Delay(usec) : 20
Avg Frame Delay(usec) : 19
Min Inter FD Variation(usec): 0
Max Inter FD Variation(usec): 1
Avg Inter FD Variation(usec): 0
FRAME DELAY BINS
Bin Number       Bin Threshold(usec)        Bin Counter
=======================================================
1               0 - < 4999                    52
2               5000 - < 9999                 0
3               10000 - < 4294967295          0
INTER-FRAME DELAY BINS
Bin Number       Bin Threshold(usec)       Bin Counter
=======================================================
1                0 - < 4999                   51
2                5000 - < 4294967295          0
```

Verify the Synthetic Loss Measurement:

```
PE1#loss-measurement type proactive profile-name SLM rmep 10 mep 20 domain 12346 ma 124
PE1#2024 Apr 10 13:40:15.587 : PE1 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame Loss
Measurement  started for MEP:20 MA:124 MD:12346
PE1#show ethernet cfm loss-measurement mep 20 domain 12346 ma-name 124
MEP: 20 MA: 124
```

```
CURRENT:
 Measurement ID : 2
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 17
  Start Time                 : 2024 Apr 10 13:41:15
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
HISTORY:
 Measurement ID : 1
  Suspect                    : False
  Measurement Type           : slm
  Elapsed time(sec)          : 60
  End Time                   : 2024 Apr 10 13:41:15
  Near End loss              : 0
  Far End loss               : 0
  Near End accumulated loss  : 0
  Far End accumulated loss   : 0
  Near End frame loss ratio  : 0
  Far End frame loss ratio   : 0
  Near End frame loss ratio min  : 0
  Far End frame loss ratio min  : 0
  Near End frame loss ratio max  : 0
  Far End frame loss ratio max  : 0
```

Implementation Examples

Enterprise Connectivity Monitoring:

Scenario: A large enterprise operates multiple branch offices connected via Ethernet services provided by a service provider network.

Use Case: Y.1731 CFM over sub-interface using Single Home EVPN ELAN enables the enterprise to monitor theconnectivity and performance of its branch office connections. It facilitates proactive fault detection and management, ensuring reliable and uninterrupted communication between the headquarters and branch offices.

Service Provider Network Operations:

Scenario: A service provider manages a diverse range of Ethernet services for its enterprise customers, including VPNs, Internet access, and cloud connectivity.

Use Case: Y.1731 CFM over sub-interface using Single Home EVPN ELAN empowers the service provider to deliverhigh-quality Ethernet services with enhanced fault management capabilities. It enables the provider to quickly identify and resolve connectivity issues, minimize service downtime, and maintain customer satisfaction.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Y.1731	A standard defined by the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) that specifies performance monitoring and fault management for Ethernet-based networks.
Sub-interface	A logical division of a physical interface, typically used to separate traffic based on VLANs or other criteria. In this context, sub-interfaces are employed to establish distinct connections within the EVPN ELAN SH topology.
EVPN	Ethernet Virtual Private Network (VPN) is a technology that enables the creation of virtual private networks over an Ethernet-based infrastructure. It provides multi-tenancy and allows for the segmentation of traffic in service provider networks.
ELAN	ELAN is a type of EVPN service that provides point-to-multi point Ethernet connectivity between two sites.
Single Home (SH)	Refers to the configuration where a Customer Edge device (CE) is connected to only one Provider Edge device (PE) within an EVPN setup. It contrasts with the multi-homed configuration, where a CE may be connected to multiple PEs.
Maintenance End Point (MEP)	MEP is a CFM entity that resides at the edge of a CFM domain. It is responsible for generating and transmitting CFM protocol packets to detect faults and collect performance data.
Maintenance Domain (MD)	MD is a logical grouping of MEPs within a CFM network. MEPs within the same MD can communicate with each other to perform CFM functions such as fault detection and performance monitoring.
Maintenance Association(MA)	MA is a collection of MEPs associated with a specific service or set of services. It defines the scope of CFM operations within a maintenance domain.
Maintenance Point Identifier (MPID)	MPID is a unique identifier assigned to each MEP within a maintenance association. It is used to distinguish between different MEPs within the same MA.
Service Level Measurement (SLM)	SLM is a CFM function used to measure the loss characteristics of a network path. It collects data on packet loss, delay, and jitter to assess the quality of service provided by the network.
Loopback Message Generation (LMM)	LMM is a CFM function used to test end-to-end connectivity by generating loopback messages. These messages are transmitted from a MEP and looped back to the same MEP to verify bidirectional communication.
Delay Measurement Message (DMM)	DMM is a CFM function used to measure the one-way delay of packets transmitted across a network. It helps assess the performance of the network in terms of packet delivery time.
Continuity Check (CC)	CC is a CFM function used to verify the continuity of a service or network path by periodically sending continuity check messages between MEPs. It helps detect connectivity faults such as link failures or misconfigurations.

CHAPTER 10 Y.1731 and CFM Over EVPN-ELAN Multi-home

Overview

The Multi Home EVPN ELAN Y.1731 CFM over Sub-interface feature enables the monitoring and management of Ethernet Virtual Private Network (EVPN) Ethernet-LAN services using the Y.1731 Connectivity Fault Management (CFM) protocol over sub-interfaces. This feature enhances fault detection and performance monitoring capabilities for EVPN E-LAN services, allowing network operators to ensure high availability and reliability of their networks. By extending Y.1731 CFM functionality to sub-interfaces in single home EVPN E-LAN deployments, this feature provides comprehensive end-to-end visibility and control, enabling proactive fault detection, isolation, and troubleshooting.

CFM multi-homing allows Customer Edge (CE) device to connect more than one Provider Edge (PE) device. Multi-homing ensures redundant connectivity. The redundant PE device ensures that there is no traffic disruption when there is a network failure.

Feature Characteristics

- Utilizes sub-interfaces to partition Ethernet traffic within the Multi Home EVPN ELAN architecture, enabling efficient service delivery and management.
- Implements EVPN ELAN architecture with single-homing capabilities, facilitating the creation of Ethernet Virtual Private Networks with simplified configurations and reduced complexity.
- Provides robust fault detection mechanisms to identify connectivity issues, link failures, and service disruptions in Ethernet networks.

Benefits

- Provides detailed insights into Ethernet service performance, enabling proactive monitoring and optimization of network resources.
- Minimizes service downtime by promptly detecting and resolving faults, ensuring uninterrupted service delivery and customer satisfaction.
- Optimizes network resource utilization and bandwidth allocation by identifying and addressing connectivity issues in a timely manner.
- Facilitates rapid fault identification and isolation, accelerating troubleshooting processes and reducing mean time to repair (MTTR).

Ensures compliance with Service Level Agreements (SLAs) by maintaining service quality metrics within defined thresholds and objectives.

Configuration

Configure Multi Home EVPN ELAN Y.1731 CFM over Sub-interface for enhanced fault management in EVPN networks.

Topology

The following topology consists of Customer Edge routers CE1 and CE2 with IPv2 Provider Edge routers PE1, PE2, and PE3. These are interconnected through the core router P in the IPv4 MPLS provider networks.

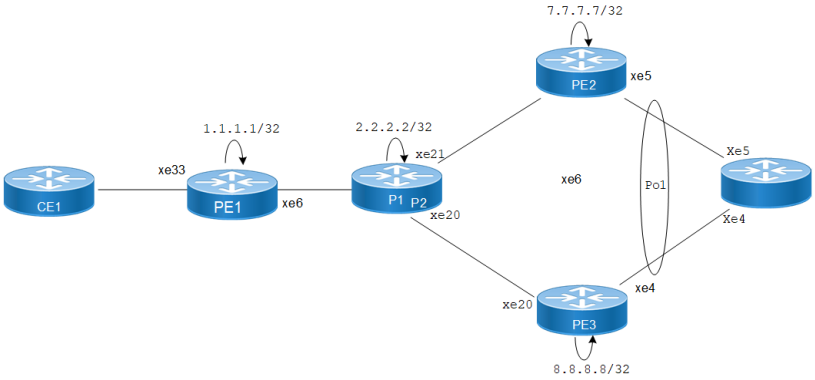

Figure 10-8: EVPN ELAN Over CFM Sub-interface

The following sessions displays the detailed information about configurations, and validations for CFM over sub-interface.

1. Configure Loopback Interface for router identification and BGP peering.

    1. Enter global configuration mode, create the loopback interface.

    ```
    PE1#configure terminal
    PE1#interface lo
    ```

    2. Assign an IP address to the loopback interface, exit interface configuration mode, and commit the changes.

    ```
    PE1(config)# interface lo
    PE1(config-if)# ip address 1.1.1.1/32
    PE1(config-if)# exit
    PE1(config)# commit
    ```

2. Configure Global LDP for distributing MPLS labels in the network.

    1. Enter LDP configuration mode.

    2. Set Router ID and configure targeted peers.

    ```
    PE1(config)# router ldp
    PE1(config-router)# router-id 1.1.1.1
    PE1(config-router)# targeted-peer ipv4 7.7.7.7
    PE1(config-router)# targeted-peer ipv4 8.8.8.8
    PE1(config-router-targeted-peer)#exit
    PE1(config-router)# exit
    PE1(config)# commit
    ```

3. Enable EVPN over MPLS and set a global VTEP IP.

    ```
    PE1(config)# evpn mpls enable
    PE1(config)# commit
    PE1(config)# evpn mpls vtep-ip-global 1.1.1.1
    PE1(config)# commit
    ```

4. Configure the interfaces connecting to the network, enabling LDP and MPLS label switching.

    ```
    PE1(config)# interface xe33
    PE1(config-if)# ip address 10.1.0.1/16
    PE1(config-if)# enable-ldp ipv4
    PE1(config-if)# label-switching
    PE1(config-if)# exit
    ```

```
PE1(config)# commit
```

5. Set up OSPF for IP routing within the network.

```
PE1(config)# router ospf 1
PE1(config-router)# ospf router-id 1.1.1.1
PE1(config-router)# network 1.1.1.1/32 area 0.0.0.0
PE1(config-router)# network 10.1.0.0/16 area 0.0.0.0
PE1(config-router)# exit
PE1(config)# commit
```

6. Set up BGP for EVPN to exchange MAC and IP information.

```
PE1(config)# router bgp 1
PE1(config-router)# neighbor 7.7.7.7 remote-as 1
PE1(config-router)# neighbor 7.7.7.7 update-source lo
PE1(config-router)# neighbor 8.8.8.8 remote-as 1
PE1(config-router)# neighbor 8.8.8.8 update-source lo
PE1(config-router)# address-family l2vpn evpn
PE1(config-router-af)# neighbor 7.7.7.7 activate
PE1(config-router-af)# neighbor 8.8.8.8 activate
PE1(config-router-af)# exit
PE1(config-router)# exit
PE1(config)# commit
```

7. Configure MAC VRF.

```
PE1(config)# mac vrf vrf2
PE1(config-vrf)# rd 1.1.1.1:2
PE1(config-vrf)# route-target both 2:2
PE1(config-vrf)# exit
PE1(config)# commit
```

8. Configure EVPN and map VRF.

```
PE1(config)# evpn mpls id 101
PE1(config-evpn-mpls)# host-reachability-protocol evpn-bgp vrf2
PE1(config)# exit
PE1(config)# commit
```

9. Configure access port on interface xe33.

```
PE1(config)# interface xe33
PE1(config-if)# interface xe33.2 switchport
PE1(config-if)# description access-side-int
PE1(config-if)# encapsulation dot1q 2
PE1(config-if)# access-if-evpn
PE1(config-access-if)# map vpn-id 101
PE1(config-access-if)# exit
PE1(config)# commit
```

10. Configure Y1731 SLM and DM profile.

```
PE1(config)#ethernet cfm loss-measurement profile-name SLM
PE1(config-cfm-lm)# measurement-type slm
PE1(config-cfm-lm)# measurement-interval 1
PE1(config-cfm-lm)# intervals-stored 3
PE1(config-cfm-lm)# message-period 1s
PE1(config-cfm-lm)# exit
PE1(config)# commit
PE1(config-cfm-lm)#ethernet cfm delay-measurement profile-name DM
PE1(config-cfm-dm)# measurement-interval 1
PE1(config-cfm-dm)# intervals-stored 2
PE1(config-cfm-dm)# message-period 1s
```

```
PE1(config-cfm-dm)#exit
PE1(config)# commit
```

Note: Similarly follow the same steps to configure respective cfm mep up and other CFM features for PE2 and PE3.

Configuration Snapshot:

PE1:

```
!
interface lo
 ip address 1.1.1.1/32
!
router ldp
 router-id 1.1.1.1
 targeted-peer 7.7.7.7
 targeted-peer 8.8.8.8
!
router ospf 1
 router-id 1.1.1.1
 network  1.1.1.1/32 area 0
 network  10.1.0.0/16 area 0
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 7.7.7.7 remote-as 1
 neighbor 7.7.7.7 update-source lo
 neighbor 8.8.8.8 remote-as 1
 neighbor 8.8.8.8 update-source lo
 !
 address-family l2vpn evpn
  neighbor 7.7.7.7 activate
  neighbor 8.8.8.8 activate
 exit-address-family
!
evpn mpls enable
evpn mpls vtep-ip-global 1.1.1.1
hardware-profile filter cfm-domain-name-str enable
hardware-profile statistics cfm-ccm enable
!
evpn mpls id 101
 host-reachability-protocol evpn-bgp vrf2

interface xe33
 ip address 10.1.0.1/16
enable-ldp ipv4
label-switching
!
vrf definition vrf2
 rd 1.1.1.1:2
 route-target both 2:2
!
evpn mpls id 52 xconnect target-mpls-id 2
 host-reachability-protocol evpn-bgp vrf2
!
interface xe33.2
 description access-side-int
 encapsulation dot1q 2
```

```
 access-if-evpn
  map vpn-id 101
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
  ethernet cfm mep up mpid 10 active true xe33.2 vlan 2
   cc multicast state enable
   exit-ether-ma- mode
   mep auto-discovery enable
   cc interval 10ms
   exit-ether-ma- mode
!
```

P:

```
!
interface lo
 ip address 2.2.2.2/32
!
interface xe6
 ip address 10.1.0.2/16
 mpls ip
!
interface xe21
 ip address 123.1.1.1/24
 enable-ldp ipv4
label-switching
!
interface xe20
 ip address 124.1.1.1/24
 enable-ldp ipv4
label-switching
!
router ldp
 router-id 2.2.2.2
!
router ospf 1
 router-id 2.2.2.2
 network 2.2.2.2/32 area 0
 network 10.1.0.0/16 area 0
 network 123.1.1.0/24 area 0
 network 124.1.1.0/24 area 0
!
```

PE2:

```
!
interface lo
 ip address 7.7.7.7/32
!
interface xe21
 ip address 123.1.1.2
 enable-ldp ipv4
label-switching
!
router ldp
 router-id 7.7.7.7/32
 targeted-peer ipv4 1.1.1.1
```

```
 targeted-peer ipv4 8.8.8.8
!
router ospf 1
 router-id 7.7.7.7
 network  7.7.7.7/32 area 0
 network 123.1.1.0/24 area 0
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 update-source lo
 neighbor 8.8.8.8 remote-as 1
 neighbor 8.8.8.8 update-source lo
 address-family l2vpn evpn
  neighbor 1.1.1.1 activate
  neighbor 8.8.8.8 activate
 exit-address-family
!
evpn mpls enable
evpn mpls vtep-ip-global 7.7.7.7
hardware-profile filter evpn-mpls-mh enable
evpn mpls multihoming enable
!
vrf definition vrf2
 rd 7.7.7.7:2
 route-target both 2:2
!
evpn mpls id 101
 host-reachability-protocol evpn-bgp vrf2
!
interface Po1
load-interval 30
evpn multi-homed system-mac 0000.aaaa.bbbc
!
interface Po1.2
switchport
encapsulation dot1q 2
access-if-evpn
map vpn-id 101
!
interface xe5
 channel-group 1 mode active
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
  ethernet cfm mep up mpid 20 active true po1.2 vlan 2
   cc multicast state enable
   exit-ether-ma- mode
   mep auto-discovery enable
   cc interval 10ms
   exit-ether-ma- mode
!
```

PE3:

```
!
interface lo
```

```
 ip address 8.8.8.8/32
!
interface xe5
 ip address 124.1.1.2/24
 enable-ldp ipv4
label-switching
!
interface xe4
 channel-group 1 mode active
!
router ldp
 router-id 8.8.8.8
 targeted-peer ipv4 1.1.1.1
 targeted-peer ipv4 7.7.7.7
!
router ospf 1
 router-id 8.8.8.8
 network  8.8.8.8/32 area 0
 network network 124.1.1.0/24 area 0
!
router bgp 1
 bgp log-neighbor-changes
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 update-source lo
 neighbor 7.7.7.7 remote-as 1
 neighbor 7.7.7.7 update-source lo
 address-family l2vpn evpn
  neighbor 1.1.1.1 activate
  neighbor 7.7.7.7 activate
 exit-address-family
!
evpn mpls enable
evpn mpls vtep-ip-global 8.8.8.8
hardware-profile filter evpn-mpls-mh enable
evpn mpls multihoming enable
!
vrf definition vrf2
 rd 8.8.8.8:2
 route-target both 2:2

evpn mpls id 101
 host-reachability-protocol evpn-bgp vrf2
!
interface Po1
load-interval 30
evpn multi-homed system-mac 0000.aaaa.bbbc
!
interface Po1.2
switchport
encapsulation dot1q 2
access-if-evpn
map vpn-id 101
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mip-
creation none
 service ma-type string ma-name 124
  ethernet cfm mep up mpid 30 active true po1.2 vlan 2
```

```
        cc multicast state enable
        exit-ether-ma- mode
        mep auto-discovery enable
        cc interval 10ms
        exit-ether-ma- mode
    !
```

Validation

The following are the validations for PE1.

PE1

The following validation is for PE1.

```
PE1#show ethernet cfm errors
domain
12346
Domain Name          Level         MEPID         Defects
12346                    7         20                .....

PE1-7011#show ethernet cfm maintenance-points remote domain 12346 ma-name 124
MEPID         RMEPID   LEVEL   Rx CCM  RDI    PEER-MAC         TYPE
10               20     7      Yes    False 00aa.bb00.0002  Learnt
10               30     7      Yes  False   00aa.dd00.0003  Learnt


PE1-7011#show ethernet cfm maintenance-points local mep domain 12346 ma-name 124 MPID
Dir Lvl CC-Stat HW-Status CC-Intvl MAC-AddressDef Port MD Name

10  Up 7Enable Installed 100 ms3417.ebe4.af22 Fxe33.2 12346

PE1-7011#ping ethernet mac 00aa.bb00.0002 unicast source 10 domain 12346 ma 124 success
rate is 100 (5/5)



PE1-7011#traceroute ethernet 00aa.bb00.0002 mepid 10 domain 12346 ma 124
MP MacHops Relay-actionIngress/Egress Ingress/Egress action 00aa.bb00.00021RlyHit
IngressIngOK

PE1-7011#ping ethernet mac 00aa.dd00.0003 unicast source 10 domain 12346 ma 124 success
rate is 100 (5/5)
PE1-7011#traceroute ethernet 00aa.dd00.0003 mepid 10 domain 12346 ma 124
MP MacHops Relay-actionIngress/Egress Ingress/Egress action 00aa.dd00.00031RlyHit
IngressIngOK
```

Verify Synthetic Loss Measurement

```
PE1#loss-measurement type proactive profile-name SLM rmep 10 mep 20 domain 12346 ma 124
PE1#2023 Sep 30 07:07:57.166 : PE1 : ONMD : INFO : [CFM_DEFECT_INFO_5]: CFM Frame Loss
Measurement started for MEP:20 MA:124 MD:12346
PE1#show ethernet cfm loss-measurement mep 20 domain 12346 ma-name 124
```

```
MEP: 20 MA: 124
CURRENT:
 Measurement ID : 2
 Suspect : False
 Measurement Type : slm
 Elapsed time(sec) : 10
 Start Time : 2023 Sep 30 07:08:56
 Near End loss : 0
 Far End loss : 0
 Near End accumulated loss : 0
 Far End accumulated loss : 0
 Near End frame loss ratio : 0
 Far End frame loss ratio : 0
HISTORY:
 Measurement ID : 1
 Suspect : False
 Measurement Type : slm
 Elapsed time(sec) : 60
 End Time : 2023 Sep 30 07:08:56
 Near End loss : 0
 Far End loss : 0
 Near End accumulated loss : 0
 Far End accumulated loss : 0
 Near End frame loss ratio : 0
 Far End frame loss ratio : 0
 Near End frame loss ratio min : 0
 Far End frame loss ratio min : 0
 Near End frame loss ratio max : 0
 Far End frame loss ratio max : 0
```

Verify Delay-measurement

```
PE1#delay-measurement type proactive profile-name DM rmep 10 mep 20 domain 12346 ma 124
PE1#2023 Oct 12 04:11:56.696 : PE1 : ONMD : INFO : [CFM_PM_SESSION_INFO_5]: CFM Frame
Delay Measurement session started for MEP Id 20 and RMEP Id 10
PE1#show ethernet cfm delay-measurement mep 20 domain 12346 ma-name 124
 MD : 12346
 MA : 124
 MEP : 20
 VC Name : test3
 Peer MAC Address : e8c5.7ae3.37ee
CURRENT:
 RMEP ID : 10
 Measurement ID : 1
 Measurement Type : DMM
 Elapsed time(sec) : 2
 Start Time : 2023 Oct 12 04:11:56
 Suspect Flag : FALSE
 Min Frame Delay(usec) : 40
 Max Frame Delay(usec) : 74
 Avg Frame Delay(usec) : 57
```

```
 Min Inter FD Variation(usec): 34
 Max Inter FD Variation(usec): 34
 Avg Inter FD Variation(usec): 34
 FRAME DELAY BINS
Bin Number Bin Threshold(usec) Bin Counter
 1 0 - < 4999 2
 2 5000 - < 9999 0
 3 10000 - < 14999 0
 4 15000 - < 4294967295 0
 INTER-FRAME DELAY BINS
Bin Number Bin Threshold(usec) Bin Counter
 1 0 - < 4999 1
 2 5000 - < 9999 0
 3 10000 - < 4294967295 0
```

PE2/PE3

The following validations for PE2 and PE3.

The following validations for PE2 and PE3.

```
PE2#show evpn mpls
EVPN-MPLS Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VPN-ID   EVI-Name     EVI-Type Type Interface ESI                          VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
101      ----         L2       NW   ----      ----                          ----      -
---      7.7.7.7          1.1.1.1
101      ----         L2       NW   ----      ----                          ----      -
---      7.7.7.7          8.8.8.8
101      ----         --       AC   po1.2 00:00:00:aa:aa:bb:bb:00:00:00 ----    DF
----             ----


Total number of entries are 4

Note: Refer sub-interface config for VLAN information.

PE3#

PE2#sh evpn mpls
EVPN-MPLS Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged
```

```
VPN-ID   EVI-Name     EVI-Type Type Interface ESI                              VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________________
_____________________________________
101      ----         L2       NW   ----      ----                             ----      -
---      8.8.8.8          1.1.1.1
101      ----         L2       NW   ----      ----                             ----      -
---      8.8.8.8          7.7.7.7
101      ----         --       AC   po1.2 00:00:00:aa:aa:bb:bb:00:00:00  ----      NON-
DF        ----             ----


Total number of entries are 4

Note: Refer sub-interface config for VLAN information.
PE3#

PE2#sh ethernet cfm errors domain 12346

Domain NameLevelMEPIDDefects

123467 20  .....
PE2#show ethernet cfm maintenance-points local mep domain 12346 ma-name 124 MPID Dir Lvl
CC-Stat HW-Status CC-Intvl MAC-AddressDef Port MD Name

20  Up 7Enable Installed 100 ms00aa.bb00.0002 Fpo1.2 12346

PE2#show ethernet cfm maintenance-points remote domain 12346 ma-name 124
MEPIDRMEPIDLEVELRx CCMRDIPEER-MACTYPE
20  10  7   Yes False 3417.ebe4.af22 Learnt PE2#ping ethernet mac 3417.ebe4.af22 unicast
source 10 domain 12346 ma 124
success rate is 100 (5/5)
PE2#traceroute ethernet 3417.ebe4.af22 mepid 10 domain 12346 ma 124
MP MacHops Relay-actionIngress/Egress Ingress/Egress action 3417.ebe4.af221RlyHit
IngressIngOK
```

CHAPTER 11 Y.1731 and CFM Over VPLS Sub-Interface

Overview

Y.1731 Connectivity Fault Management (CFM) over Layer 2 Virtual Private LAN Service (VPLS) is a protocol and technology combination used for fault management in Layer 2 VPN networks. It allows for the detection and management of faults, performance monitoring, and fault localization within a VPLS network

Feature Characteristics

- Facilitates end-to-end fault management across the VPLS network, covering provider and customer edges.
- Supports multi-level fault management, allowing operators to define different levels of fault detection and management for different parts of the network.
- Y.1731 CFM includes performance monitoring capabilities, such as delay measurement and frame loss measurement, to monitor service quality parameters.
- The protocol supports loopback and link trace functions to identify and troubleshoot faults within the VPLS network.

Benefits

- Enables rapid detection and localization of faults within the VPLS network, minimizing downtime and service disruptions.
- Provides performance monitoring capabilities, allowing to track key performance indicators and ensure service quality.
- Enhances network visibility by providing detailed fault and performance monitoring data, aiding in network troubleshooting and maintenance.

Prerequisites

Ensure the network devices participating in the L2VPN VPLS setup support Y.1731 CFM functionality. This includes the Provider Edge (PE) and Customer Edge (CE) devices.

Configuration

Configure Y.1731 CFM over sub-interface using L2VPN VPLS by defining the CFM domain, configuring service MEPs and MAs, and setting up cross-connects between primary and backup interfaces.

Topology

The topology consists of three Customer Edge devices (CE1, CE2, and CE3) connected to three Provider Edge devices (PE1, PE2, and PE3) via sub-interfaces (xe1, xe12, and xe13). The Provider Edge devices are interconnected through Provider Devices (P1 and P2). Y.1731 ethernet CFM is configured over these sub-interfaces to monitor and manage ethernet connectivity between the CE devices, ensuring fault detection and performance monitoring across the service provider's network.

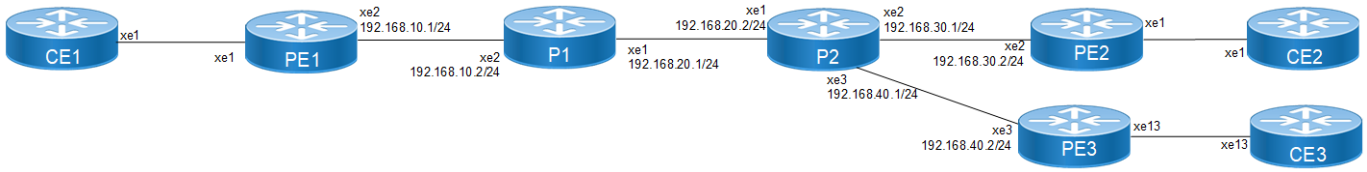

Figure 11-9: L2VPN VPLS Y1731 CFM Over Sub-interface

Perform the following configurations to configure Y.1731 CFM over sub-interface using L2VPN VPLS:

1. On Customer Edge (CE) Nodes (CE1, CE2, and CE3), configure the interface xe1 and set it as a switchport with a load interval of (30 seconds):

```
CE1(config)#interface xe1
CE1(config-if)#switchport
CE1(config-if)#load-interval 30
CE1(config-if)#commit
CE1(config-if)#exit
```

Note: Similarly follow the same steps to configure xe11(CE1), xe12(CE2), and xe13(CE3).

2. Create sub-interface (xe1.2001)adding the VLAN:

```
CE1(config)#interface xe1.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit
CE1(config)#interface xe11.2001 switchport
CE1(config-if)#encapsulation dot1q 2028
CE1(config-if)#commit
CE1(config-if)#exit
```

3. Set up a cross-connect named (`test100`), specifying in and out interfaces:

```
CE1(config)#cross-connect test100
CE1(config-xc)#interface xe1.2001
CE1(config-xc)#interface xe11.2001
CE1(config-xc)#commit
```

4. Perform the following on PE1:

    1. Configure CFM related hardware profiles:

```
PE1(config)# hardware-profile filter cfm-domain-name-str enable
PE1(config)# hardware-profile statistics cfm-lm enable
PE1(config)# hardware-profile statistics cfm-ccm enable
PE1(config)#hardware-profile statistics cfm-slm enable
```

    2. Configure the loopback interface with a secondary IP address(`1.1.1.1/32`):

```
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#commit
PE1(config-if)#exit
```

    3. Configure LDP targeted peers:

```
PE1(config)#router ldp
PE1(config-router)#targeted-peer ipv4 4.4.4.4
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#transport-address ipv4 1.1.1.1
```

```
PE1(config-router)#commit
PE1(config-router)exit
```

4. Configure interface `xe2` with an IP address (`192.168.10.1/24`) and enable LDP:

```
PE1(config)#interface xe2
PE1(config-if)#load-interval 30
PE1(config-if)#ip address 192.168.10.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#commit
PE1(config-if)#exit
```

5. Configure OSPF routing, specify the OSPF router ID as (`1.1.1.1`), enable BFD on all interfaces, define the network (`1.1.1.1/32`) in area (`0.0.0.0`), and define the network (`192.168.10.0/24`) in area (`0.0.0.0`):

```
PE1(config)#router ospf 1
PE1(config-router)#ospf router-id 1.1.1.1
PE1(config-router)#bfd all-interfaces
PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0
PE1(config-router)#network 192.168.10.0/24 area 0.0.0.0
PE1(config-router)#commit
PE1(config-router)#exit
```

6. Set up an L2VPN VPLS between PE1, PE2, and PE3.

```
PE1(config)#mpls vpls vpls-301 301
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)#  vpls-type vlan
PE1(config-vpls-sig)#  vpls-peer 4.4.4.4
PE1(config-vpls-sig)#  vpls-peer 5.5.5.5
PE1(config-vpls-sig)#  exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#commit
PE1(config)#exit
```

7. Configure sub-interface (xe1.2001) as an access interface for VPLS.

```
PE1(config)#interface xe1.2001 switchport
PE1(config-if)#encapsulation dot1q 2028
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)#mpls-vpls vpls-301
PE1(config-acc-if-vpls)#commit
PE1(config-acc-if-vpls)#exit
```

8. Configure Up-mep CFM domain:
   - Set the domain type as a character string with the domain name (12346) and (level 7)
   - Specify the MA type as a string with the MA name (124)
   - Set up a MEP with MEP ID (20) as active on interface (xe1.2001) and Associate the vlan (VLAN 2028)
   - Enable multicast state for continuity check, and auto-discovery of RMEPs
   - Set the continuity check interval to (10 milliseconds)

```
PE1(config)#ethernet cfm domain-type character-string domain-name
12346 level 7 mip-creation none
PE1(config-ether-cfm)# service ma-type string ma-name 124
PE1(config-ether-cfm-ma)#ethernet cfm mep up mpid 20 active true
xe1.2001 vlan 2028
PE1(config-ether-cfm-ma-mep)#cc multicast state enable
PE1(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode
```

```
PE1(config-ether-cfm-ma)#rmep auto-discovery enable
PE1(config-ether-cfm-ma)#cc interval 10ms
PE1(config-ether-cfm-ma)#exit-ether-ma-mode
PE1(config-ether-cfm)#commit
PE1(config-ether-cfm)exit
```

- Create a loss measurement profile named SLM with measurement type SLM, measurement interval of 1, intervals stored of 3, and message period of (1) second.

```
PE1(config)#ethernet cfm loss-measurement profile-name SLM
PE1(config-cfm-lm)#measurement-type slm
PE1(config-cfm-lm)#measurement-interval 1
PE1(config-cfm-lm)#intervals-stored 3
PE1(config-cfm-lm)#message-period 1s
PE1(config-cfm-lm)#exit
```

- Create a delay measurement profile named DM with a measurement interval of (1), intervals stored of (2), and message period of (1 second).

```
PE1(config)#ethernet cfm delay-measurement profile-name DM
PE1(config-cfm-dm)#measurement-interval 1
PE1(config-cfm-dm)#intervals-stored 2
PE1(config-cfm-dm)#message-period 1
```

Configuration Snapshot:

PE1:

```
interface lo
 ip address 1.1.1.1/32 secondary
!
router ldp
 targeted-peer ipv4 4.4.4.4
  exit-targeted-peer-mode
 targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
transport-address ipv4 1.1.1.1
!
mpls vpls vpls-301 301
 signaling ldp
  vpls-type vlan
  vpls-peer 4.4.4.4
  vpls-peer 5.5.5.5
  exit-signaling
 exit-vpls
!
interface xe2
 load-interval 30
 ip address 192.168.10.1/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 192.168.10.0/24 area 0.0.0.0
```

```
!
interface xe1
 switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-vpls
  mpls-vpls vpls-301
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mipcreation none
 service ma-type string ma-name 124
 ethernet cfm mep up mpid 20 active true xe1.2001 vlan 2028
 cc multicast state enable
 exit-ether-ma-mep-mode
 rmep auto-discovery enable
 cc interval 10ms
 exit-ether-ma-mode
!
ethernet cfm loss-measurement profile-name SLM
 measurement-type slm
 measurement-interval 1
 intervals-stored 3
 message-period 1s
!
ethernet cfm delay-measurement profile-name DM
 measurement-interval 1
 intervals-stored 2
 message-period 1s
```

PE2:

```
interface lo
 ip address 4.4.4.4/32 secondary
!
router ldp
 targeted-peer ipv4 1.1.1.1
exit-targeted-peer-mode
targeted-peer ipv4 5.5.5.5
exit-targeted-peer-mode
transport-address ipv4 4.4.4.4
!
interface xe2
 load-interval 30
 ip address 192.168.30.2/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id 4.4.4.4
 bfd all-interfaces
```

```
 network 4.4.4.4/32 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0
!
mpls vpls vpls-301 301
 signaling ldp
  vpls-type vlan
  vpls-peer 1.1.1.1
  vpls-peer 5.5.5.5
  exit-signaling
 exit-vpls
!
interface xe1
 switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-vpls
  mpls-vpls vpls-301
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mipcreation none
 service ma-type string ma-name 124
 ethernet cfm mep up mpid 10 active true xe1.2001 vlan 2028
 cc multicast state enable
 ethernet cfm loss-measurement reply slm
 ethernet cfm delay-measurement reply dmm
 exit-ether-ma-mep-mode
 rmep auto-discovery enable
 cc interval 10ms
 exit-ether-ma-mode
!
```

PE3:

```
interface lo
 ip address 5.5.5.5/32 secondary
!
router ldp
 targeted-peer ipv4 1.1.1.1
exit-targeted-peer-mode
targeted-peer ipv4 4.4.4.4
exit-targeted-peer-mode
transport-address ipv4 5.5.5.5
!
interface xe3
 load-interval 30
 ip address 192.168.40.2/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id 5.5.5.5
```

```
 bfd all-interfaces
 network 5.5.5.5/32 area 0.0.0.0
 network 192.168.40.0/24 area 0.0.0.0
!
mpls vpls vpls-301 301
 signaling ldp
  vpls-type vlan
  vpls-peer 1.1.1.1
  vpls-peer 4.4.4.4
  exit-signaling
 exit-vpls
!
interface xe1
 switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
 access-if-vpls
  mpls-vpls vpls-301
!
ethernet cfm domain-type character-string domain-name 12346 level 7 mipcreation none
 service ma-type string ma-name 124
 ethernet cfm mep up mpid 30 active true xe1.2001 vlan 2028
 cc multicast state enable
  ethernet cfm loss-measurement reply slm
 ethernet cfm delay-measurement reply dmm
 exit-ether-ma-mep-mode
 rmep auto-discovery enable
 cc interval 10ms
 exit-ether-ma-mode
!
```

P2:

```
interface lo
 ip address 3.3.3.3/32 secondary
!
router ldp
 transport-address ipv4 3.3.3.3
!
interface xe1
 ip address 192.168.20.2/24
 label-switching
 enable-ldp ipv4
!
interface xe2
 ip address 192.168.30.1/24
 label-switching
 enable-ldp ipv4
!
interface xe3
```

```
 ip address 192.168.40.1/24
 label-switching
 enable-ldp ipv4
!
router ospf 1
 ospf router-id 3.3.3.3
 bfd all-interfaces
 network 3.3.3.3/32 area 0.0.0.0
 network 192.168.20.0/24 area 0.0.0.0
 network 192.168.30.0/24 area 0.0.0.0
 network 192.168.40.0/24 area 0.0.0.0
```

CE3:

```
interface xe1
switchport
 load-interval 30
!
interface xe1.2001 switchport
 encapsulation dot1q 2028
!
interface xe13.2001 switchport
 encapsulation dot1q 2028
!
cross-connect test100
 interface xe1.2001
 interface xe13.2001
```

CE1:

```
    interface xe1
    switchport
     load-interval 30

    interface xe1.2001 switchport
     encapsulation dot1q 2028

    interface xe11.2001 switchport
     encapsulation dot1q 2028

    cross-connect test100
     interface xe1.2001
     interface xe11.2001
```

CE2:

```
    interface xe1
    switchport
     load-interval 30

    interface xe1.2001 switchport
     encapsulation dot1q 2028

    interface xe12.2001 switchport
     encapsulation dot1q 2028

    cross-connect test100
```

```
     interface xe1.2001
     interface xe12.2001
```

Validation

Verify the L2VPN VPLS status.

```
=========================
PE1#  show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr         Tunnel-Label  In-Label   Network-Intf   Out-Label  Lkps/St
PW-INDEX SIG-Protocol   Status     UpTime
301         4.4.4.4           52481         26240      xe2            28160      2/Up
2          LDP            Active     1d00h02m
301         5.5.5.5           52497         26256      xe2            26256      2/Up
3          LDP            Active     1d00h57m
PE1#
```

Verify the RMEP is learned or not.

```
PE1#show ethernet cfm maintenance-points remote domain 12346
 MA_NAME MEPID RMEPID LEVEL Rx CCM RDI PEER-MAC TYPE
--------------------------------------------------------------------------------
 124 20 10 7 Yes False e8c5.7ae3.37ee Learnt
 124 20 30 7 Yes False e8c5.7ae3.38ee Learnt
```

Verify the Ping:

```
PE1#ping ethernet mac e8c5.7ae3.37ee unicast source 20 domain 12346 ma 124
 success rate is 100 (5/5)

PE1#ping ethernet mac e8c5.7ae3.38ee unicast source 20 domain 12346 ma 124
 success rate is 100 (5/5)
```

Verify the Traceroute:

```
PE1#traceroute ethernet e8c5.7ae3.37ee mepid 20 domain 12346 ma 124
 MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
e8c5.7ae3.37ee 1 RlyHit Ingress IngOK

PE1#traceroute ethernet e8c5.7ae3.38ee mepid 20 domain 12346 ma 124
 MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
e8c5.7ae3.38ee 1 RlyHit Ingress IngOK
```

Implementation Examples

- To support a vast network infrastructure delivering VPLS to a multitude of enterprise clients, it is imperative to maintain uninterrupted connectivity and peak performance for these VPLS connections, all while minimizing the risk of downtime or disruptions.
- Understanding the role of fault detection, localization, and performance monitoring within the VPLS network, deploy Y.1731 CFM over Layer 2 VPN (VPLS) to enhance the network's resilience and operational efficiency.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Virtual Private LAN Service (VPLS)	Allows multiple sites in different geographical locations to connect over a wide area network (WAN), creating the experience of being part of a single local area network (LAN).
Connectivity Fault Management (CFM)	CFM is a protocol used to detect, verify, and isolate connectivity faults in a network. It operates at the data link layer and is designed to monitor ethernet networks.
Virtual Private LAN Service (VPLS)	Allows multiple dispersed sites to connect over a wide area network (WAN), creating the experience of being part of a single local area network (LAN).
Maintenance End Point (MEP)	MEP is a CFM entity that resides at the edge of a CFM domain. It is responsible for generating and transmitting CFM protocol packets to detect faults and collect performance data.
Maintenance Domain (MD)	MD is a logical grouping of MEPs within a CFM network. MEPs within the same MD can communicate with each other to perform CFM functions such as fault detection and performance monitoring.
Maintenance Association(MA)	MA is a collection of MEPs associated with a specific service or set of services. It defines the scope of CFM operations within a maintenance domain.
Maintenance Point Identifier (MPID)	MPID is a unique identifier assigned to each MEP within a maintenance association. It is used to distinguish between different MEPs within the same MA.
Service Level Measurement (SLM)	SLM is a CFM function used to measure the loss characteristics of a network path. It collects data on packet loss, delay, and jitter to assess the quality of service provided by the network.
Loopback Message Generation (LMM)	LMM is a CFM function used to test end-to-end connectivity by generating loopback messages. These messages are transmitted from a MEP and looped back to the same MEP to verify bidirectional communication.
Delay Measurement Message (DMM)	DMM is a CFM function used to measure the one-way delay of packets transmitted across a network. It helps assess the performance of the network in terms of packet delivery time.
Continuity Check (CC)	CC is a CFM function used to verify the continuity of a service or network path by periodically sending continuity check messages between MEPs. It helps detect connectivity faults such as link failures or misconfigurations.

Improved Management

This section describes the network monitoring enhancements and new features introduced in the Release 6.5.3.

Release 6.5.3

- In-band Management over Custom VRF

Release 6.5.2

- Streaming Telemetry Dial-Out Mode
- Streaming Telemetry Dial-Out Mode
- DHCPv6 Prefix Delegation Configuration
- Configure SRv6 with EVPN ELAN
- Configure SRv6 with EVPN ELAN
- BGP ORF Prefix-List VPNV4 Address

CHAPTER 1 In-band Management over Custom VRF

Overview

OcNOS currently supports system management protocols within the Default and Management Virtual Routing and Forwarding (VRF). However, this configuration is insufficient for customer deployments that require the ability to run these protocols in user-defined VRFs. This document outlines the requirements for expanding OcNOS to support system management protocols in custom VRFs.

Feature Characteristics

- **Support for System Management Protocols in User-Defined VRFs:** Provide the flexibility to run system management protocols over custom VRFs. In large-scale networks, deploying an out-of-band management network is not always practical, making in-band device management over user-defined VRFs necessary to handle the volume of management traffic.
- **Supported Protocols:** SSH, Telnet, TACACS, Syslog, SNMP, NETCONF, and gNMI will operate within user-defined VRFs. Simultaneous support for multiple VRFs for specific protocols, such as Syslog. Support for both default and customizable port values for each protocol.
- **Multi-VRF Protocol Operations:** Management protocols, including SSH and NETCONF, will allow simultaneous operations across multiple VRFs, providing enhanced flexibility in managing network devices.
- **Service Traffic Segmentation:** Management traffic, such as SNMP and Syslog, can be segmented across custom VRFs, allowing for more efficient traffic management and security.

Benefits

- **Scalability and Flexibility:** Enabling system management protocols to operate over custom VRFs allows for ease of managing service provider networks, especially in environments where out-of-band management is impractical.
- **Protocol Customization:** Support for both standard and customizable port values for management protocols provides greater flexibility, allowing customers to tailor the system management configuration to meet their specific network needs.

Configuration

These steps provide a standardized approach to configuring User-Defined VRF on PE routers.

Topology

In this topology, the management traffic from the Linux Server is routed through a specific VRF that is isolated from the traffic on the L3VPN.

PE1 and PE2 are Provider Edge routers in the network. These routers are responsible for managing and routing the traffic between the local network and the wider service provider network.

Both PE routers are connected through L3VPN, which is used to segment and isolate traffic between the two routers over a shared infrastructure. Each customer or service can have its own isolated routing table (VRF).

In-Band Connection: The In-Band Connection shown between PE1 and the Linux Server means that both management and normal traffic flow over the same physical network links.

The in-band management traffic is directed over the custom VRF, ensuring it is separated from the service traffic, providing network isolation.

Custom VRF Feature: In this case, the custom VRF is applied to manage the traffic between the Linux Server and the network. This VRF allows traffic related to management tasks to remain separate from other traffic handled by the provider.

VRF helps ensure that different traffic types (such as syslog, or SSH sessions) remain isolated for security and performance reasons.

Multi-VRF Management: Using user-defined VRFs, run management services like Syslog, or SSH on separate VRFs, ensuring that management tasks are not mixed with customer or service traffic.

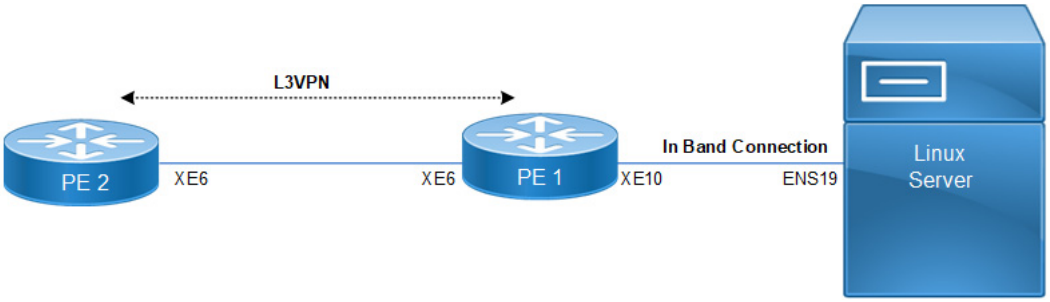

Custom VRF

Perform the following configuration steps for setting up a custom VRF with routing protocols like BGP, OSPF, and management protocols such as SSH. These can be applied to multiple Provider Edge (PE) routers, or other routers, with adjustments in interface names and IP addresses depending on the specific deployment.

The steps include defining VRFs, configuring interfaces, setting up routing protocols like OSPF and BGP, enabling management features (SSH), and ensuring MPLS support:

1. Enter configuration mode and define the VRF.

```
#configure terminal
(config)# ip vrf vrf1
(config)# rd 100:1
(config)# route-target both 10:10
(config)#exit
```

2. Assign the VRF to the relevant access and loopback interfaces, and configure both IPv4 or IPv6 addresses:

   Access Interface Configuration:

```
#configure terminal
(config)# interface xe10
(config)# ip vrf forwarding vrf1
(config)# ip address 20.20.20.3/24
(config)# ipv6 address 2500::3/64
(config)#exit
```

   Loopback Interface Configuration:

```
#configure terminal
(config)# interface lo.vrf1
(config)# ip vrf forwarding vrf1
(config)# ip address 172.16.1.10/24 secondary
(config)# ipv6 address 2000::10/64
(config)#exit
```

3. On interfaces facing the provider network, configure MPLS and enable LDP:

```
(config)# interface xe6
(config)# ip address 192.168.69.1/24
(config)# ipv6 address 1000::11/64
(config)# label-switching
(config)# enable-ldp ipv4
(config)#exit
```

4. Set up OSPF routing within the network, and ensure to advertise the necessary interfaces:

```
(config)# router ospf 100
(config)# network 1.1.1.1/32 area 0.0.0.0
(config)# network 192.168.69.0/24 area 0.0.0.0
(config)#commit
(config)#exit
#configure terminal
(config)# router ldp
(config)#exit
```

5. Configure BGP for both VPNv4 and VPNv6 address families:

```
#configure terminal
(config)# router bgp 1000
(config)# neighbor 2.2.2.2 remote-as 1000
(config)# neighbor 2.2.2.2 update-source 1.1.1.1
(config)# address-family vpnv4 unicast
(config)# neighbor 2.2.2.2 activate
(config)# exit-address-family
(config)# address-family ipv4 vrf vrf1
(config)# redistribute connected
(config)# exit-address-family
(config)# address-family vpnv6 unicast
(config)# neighbor 2.2.2.2 activate
(config)# exit-address-family
(config)# address-family ipv6 vrf vrf1
(config)# redistribute connected
(config)# exit-address-family
(config)#exit
```

6. Enable SSH (or respective protocols) for VRF Management:

```
#configure terminal
(config)# feature ssh vrf management
(config)# feature ssh vrf
(config)# feature ssh vrf vrf1
(config)#exit

(config)# ssh server port 10000 vrf management
(config)# ssh server port 10000
(config)# ssh server port 10000 vrf vrf1
(config)#exit

(config)# ssh login-attempts 2 vrf management
(config)# ssh login-attempts 2
(config)# ssh login-attempts 2 vrf vrf vrf1
(config)#exit

(config)# ssh server session-limit 10 vrf management
(config)# ssh server session-limit 10
(config)# ssh server session-limit 20 vrf vrf1
(config)#exit
```

```
(config)# ssh server algorithm encryption 3des-cbc  vrf management
(config)# ssh server algorithm encryption 3des-cbc
(config)# ssh server algorithm encryption 3des-cbc  vrf vrf1
(config)#exit
```

Configuration Snapshot:

```
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
logging console
logging monitor
logging cli
logging level all 7
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature ssh vrf vrf1
ssh server port 10000 vrf vrf1
ssh login-attempts 2 vrf vrf1
ssh server algorithm encryption 3des-cbc  vrf vrf1
ssh server session-limit 20 vrf vrf1
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
!
ip vrf management
!
```

```
ip vrf vrf1
 rd 100:1
 route-target both 10:10
!
router ldp
!
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.vrf1
 ip vrf forwarding vrf1
 ip address 172.16.1.10/24 secondary
 ipv6 address 2000::10/64
!
interface xe6
 speed 10g
 ip address 192.168.69.1/24
 ipv6 address 1000::11/64
 label-switching
 enable-ldp ipv4
!
 ip vrf forwarding vrf1
 ip address 20.20.20.3/24
 ipv6 address 2500::3/64
!
!
router ospf 100
 network 1.1.1.1/32 area 0.0.0.0
 network 192.168.69.0/24 area 0.0.0.0
!
router bgp 1000
 neighbor 2.2.2.2 remote-as 1000
 neighbor 2.2.2.2 update-source 1.1.1.1
 !
 address-family vpnv4 unicast
 neighbor 2.2.2.2 activate
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 2.2.2.2 activate
```

```
 exit-address-family
 !
 address-family ipv4 vrf vrf1
 redistribute connected
 exit-address-family
 !
 address-family ipv6 vrf vrf1
 redistribute connected
 exit-address-family
 !
 exit
!
line console 0
 exec-timeout 0 0
line vty 0
 exec-timeout 0 0
!
```

Validation

Validate the VRF and SSH configurations to ensure they support the custom VRF functions as expected.

- **Verify the VRF Configuration:**

```
OcNOS#show running-config vrf vrf1
!
ip vrf vrf1
 rd 100:1
 route-target both 10:10
!
OcNOS#show running-config interface xe10
!
interface xe10
 ip vrf forwarding vrf1
 ip address 20.20.20.3/24
 ipv6 address 2500::3/64
!
OcNOS#show running-config interface lo.vrf1
!
interface lo.vrf1
 ip vrf forwarding vrf1
 ip address 172.16.1.10/24 secondary
 ipv6 address 2000::10/64
!
OcNOS#show running-config interface xe6
interface xe6
 speed 10g
 ip address 192.168.69.1/24
 ipv6 address 1000::11/64
 label-switching
 enable-ldp ipv4
!
```

- **Verify SSH configuration:**

```
OcNOS#show running-config ssh server
feature ssh vrf management
no feature ssh
feature ssh vrf vrf1
ssh server port 10000 vrf vrf1

OcNOS#show ssh server
VRF MANAGEMENT:
ssh server enabled port: 22
authentication-retries: 3
VRF DEFAULT:
ssh server disabled port: 22
authentication-retries: 3
VRF vrf1:
ssh server enabled port: 10000
session-limit: 20
authentication-retries: 2
```

Implementation Examples

- L3VPN or EVPN Tunnel Support: In a service provider network, user-defined VRFs are configured on managed nodes, such as PE and Rout Reflector (RR) nodes. Management nodes connect to a PE node, enabling access to other PE or RR nodes through L3VPN or EVPN tunnels. This architecture supports in-band management of devices over user-defined VRFs.
- Service Traffic Segmentation: Management traffic, such as SNMP and Syslog packets, is segmented across different user-defined VRFs, ensuring separation from other network operations and enhancing security.
- Multi-VRF Support for Protocols: SSH and NETCONF services support connections from multiple VRFs simultaneously, allowing for scalable management across complex networks.

Glossary

The following list provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Virtual Routing and Forwarding (VRF)	A technology that allows multiple instances of a routing table to coexist on the same router. Each VRF operates independently, enabling isolated network paths and address spaces within a single physical infrastructure.
Multiprotocol Label Switching (MPLS)	A method for forwarding packets based on labels rather than network addresses. MPLS is commonly used in conjunction with VRF to route traffic through the network efficiently.
Label Distribution Protocol (LDP)	A protocol used in MPLS networks to establish label-switched paths (LSPs). LDP is responsible for distributing labels between routers to forward packets in an MPLS environment.

Open Shortest Path First (OSPF)	A link-state interior gateway protocol (IGP) used to distribute IP routing information within a single autonomous system. It is commonly used in conjunction with VRFs to handle routing within a VRF instance.
Border Gateway Protocol (BGP)	The protocol used to exchange routing information between different autonomous systems. When combined with VRFs, BGP can handle VPNv4 and VPNv6 routes for isolated routing domains.
Secure Shell (SSH)	A protocol that provides secure access to network devices and systems. In a VRF configuration, SSH can be enabled per VRF, allowing secure management access to routers on a per-VRF basis.

CHAPTER 2 Streaming Telemetry Dial-Out Mode

Overview

In OcNOS, dial-out telemetry subscriptions, also known as persistent subscriptions, ensure continuous data streaming, even if the Remote Procedure Call (gRPC) session terminates unexpectedly. With persistent subscriptions, the OcNOS device continuously retries to establish a gRPC connection to the collector server, thus maintaining persistent data streaming.

Feature Characteristics

The dial-out telemetry feature in OcNOS comprises several key aspects ensuring seamless data streaming and connectivity with collector servers:

The described topology outlines a system architecture that utilizes gRPC-based tunneling for persistent streaming telemetry.

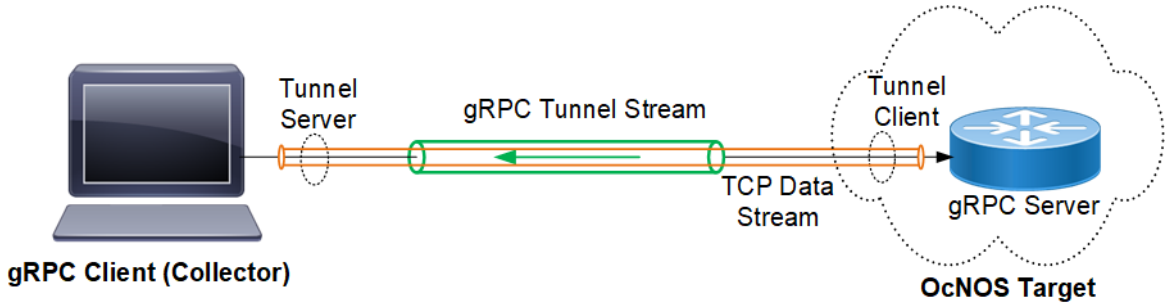

Figure 2-10: Dial-Out Subscription Mode

Here is a detailed explanation of the components and data flow:

- **gNMI Client (gRPC Client):** The gNMI client, which acts as the gRPC client in this scenario, is responsible for handling telemetry data and connecting to the OcNOS target device.
- **Tunnel Server:** The tunnel server, part of the gNMI collector process, listens for incoming gRPC tunnel streams from the gRPC server.
- **gRPC Tunnel Stream:** Represents the persistent communication channel established between the tunnel client (OcNOS) and the tunnel server (collector).
- **Tunnel Client:** The gRPC tunnel client operates on the OcNOS device and connects to the tunnel server. It manages the tunneling of telemetry data.
- **gRPC Server:** Interacts with the tunnel client to establish and manage the tunnel.

Note: Ensure that the tunnel server is reachable over the network from the tunnel client, and configure both the tunnel client and tunnel server with compatible authentication mechanisms.

Data Flow

The Data Flow: Dial-Out Mode flow chart illustrates streaming telemetry in Dial-out Mode.

- **Initialization:** When the dial-out command subscription-name is applied successfully, the tunnel client on the OcNOS device initiates a connection to the tunnel server hosted on the collector.
- **Tunnel Establishment:** Upon successful connection, the gRPC client and server establish a persistent tunnel stream. This tunnel facilitates the continuous transmission of telemetry data.

Note: OcNOS supports insecure tunnel connections.

- **Telemetry Data Transmission:** When telemetry data needs to be transmitted from the OcNOS device, the gNMI client sends a Publish RPC request over the established tunnel.
- **Subscription Configuration:** Telemetry commands follow the OpenConfig telemetry model, standardizing the configuration of telemetry subscriptions and related entities.

Figure 2-11: Data Flow: Dial-Out Mode

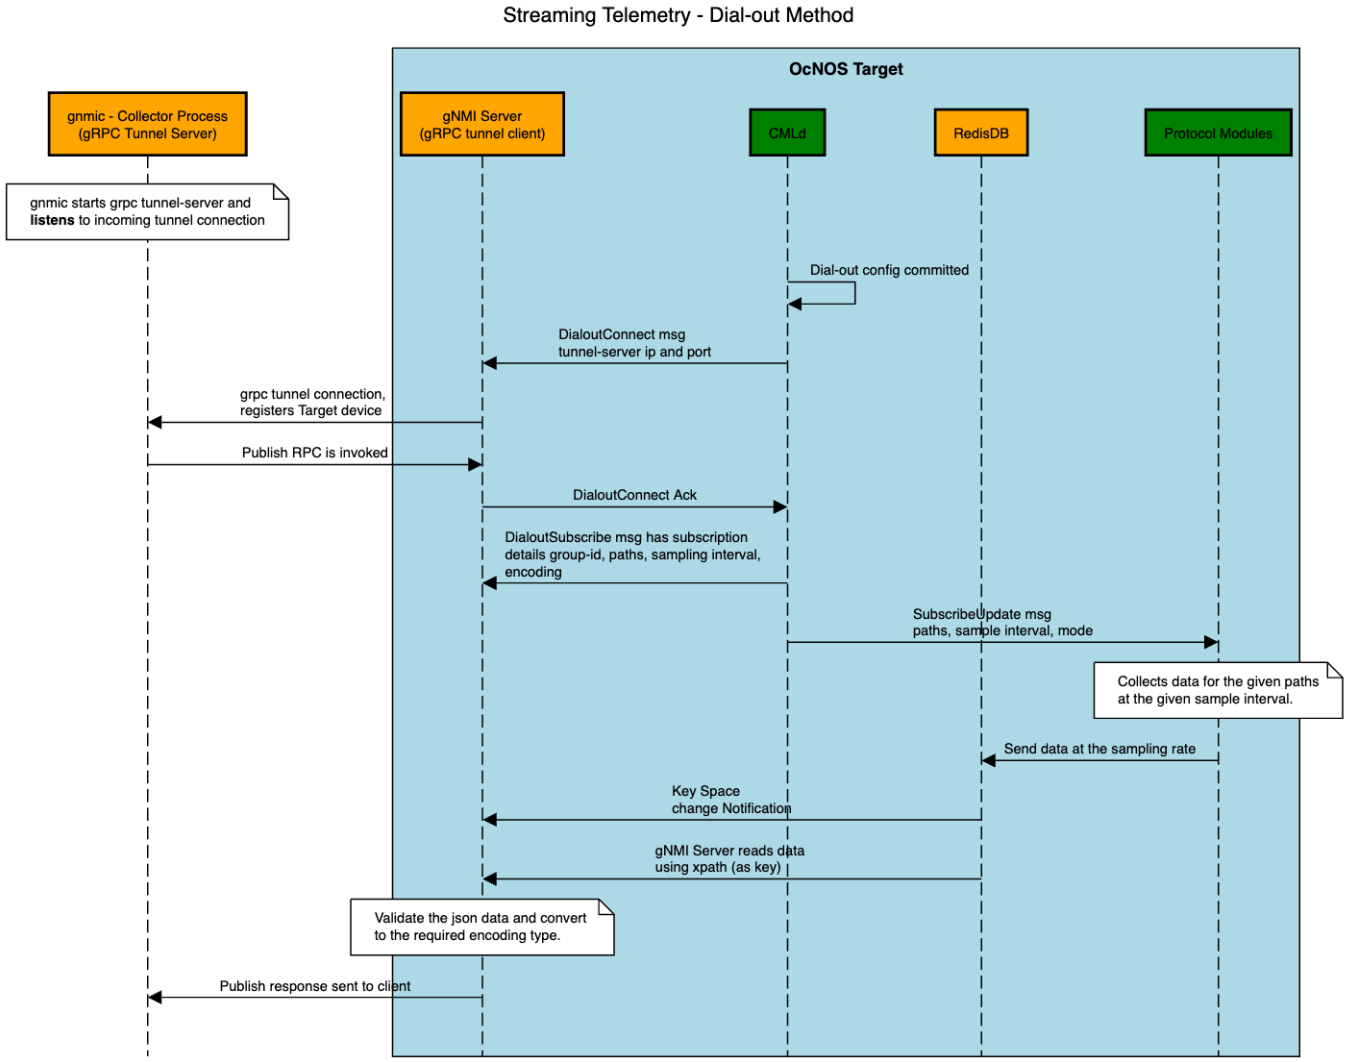

Benefits

- Ensures continuous data streaming even in the event of gRPC session termination, enhancing network monitoring and troubleshooting capabilities.
- Simplifies configuration and management of telemetry subscriptions using standard OpenConfig models.
- Facilitates secure and reliable communication between the OcNOS device and the collector server.
- Enhances interoperability by enabling integration with third-party gRPC client applications like gNMI client, expanding telemetry options for network operators.

Prerequisites

Before configuring Dial-Out mode, ensure that:

- A supported OcNOS router running a compatible release is required.
- Access to the management interface of the router is necessary.
- Refer to the gnmic Installation to download the gNMI collector package.

Configuration

Set up the OcNOS router to transmit streaming telemetry data to a gNMI client using the dial-out method.

The sample configuration on the OcNOS router sets up streaming telemetry subscriptions using gNMI to monitor specific paths related to the state of Hard Disk, RAM, and Chassis. The router sends telemetry data to the specified collector over a configured tunnel connection. The gNMI client subscribed to these paths will receive updates regarding the state of RAM and Hard Disk at the specified intervals. This setup enables proactive monitoring and management of key hardware components on the network device.

Topology

In this setup, an OcNOS router functions as the data source for streaming telemetry, while a gNMI client acts as the receiver of telemetry data. The OcNOS router sends telemetry data to the gNMI client over a dial-out connection.

Figure 2-12: Dial-out Streaming Telemetry Topology

Use Case 1: Configure Telemetry on Management VRF

Note: Before configuring Dial-out, meet all Prerequisites.

1. Enable Streaming Telemetry on a management VRF.

```
OcNOS(config)#feature streaming-telemetry vrf management
```

2. Create Sensor Group

   Create a sensor group (`Platform`) where sensor paths will be specified for dial-out subscriptions. Specify sensor paths within the sensor group (`Platform`) to monitor the chassis state.

```
OcNOS(config)#sensor-group Platform vrf management
OcNOS(telemetry-sensor-group)#sensor-path ipi:/components/
component[name=CHASSIS]/state
```

```
OcNOS(telemetry-sensor-group)#exit
```

3. Create Destination Group

   Create a destination group (`Collector2`) where tunnel server settings will be configured for dial-out subscriptions. Specify the tunnel server (`gNMI Client`) IP address (`10.21.3.4`) and port (`11123`) within the destination group (`Collector2`).

```
OcNOS(config)#destination-group Collector2 vrf management
OcNOS(telemetry-grpc-tunnel-group)#tunnel-server ip 10.21.3.4 port 11123
OcNOS(telemetry-grpc-tunnel-group)#exit
```

4. Create Persistent Subscription

   Create a persistent subscription (`storage2`), encoding type (`JSON-IETF`), and associate it with the destination group (`Collector2`), and sensor group (`Platform`) to monitor the chassis state with a sample interval (`95 seconds`).

```
OcNOS(config)#subscription-name storage2 vrf management
OcNOS(telemetry-subscription)#encoding json-ietf
OcNOS(telemetry-subscription)#destination-group Collector2
OcNOS(telemetry-subscription)#sensor-group Platform sample-interval 95
OcNOS(telemetry-subscription)#commit
OcNOS(telemetry-subscription)#exit
```

Streaming Telemetry Snippet Configurations on Management VRF

To verify the telemetry configuration and view the overall commands used for dial-out subscriptions, use the `show running-config streaming-telemetry` command on the router.

```
OcNOS#show running-config streaming-telemetry
!
feature streaming-telemetry vrf management
!
sensor-group Platform vrf management
 sensor-path ipi:/components/component[name=CHASSIS]/state
!
destination-group Collector2 vrf management
 tunnel-server ip 10.21.3.4 port 11123
!
subscription-name storage2 vrf management
 destination-group Collector2
 sensor-group Platform sample-interval 95
!
!
```

Use Case 2: Configure Telemetry on User-defined VRF

Note: Before configuring Dial-out, meet all Prerequisites.

1. Enable Streaming Telemetry in a user-defined VRF on an OcNOS router.

```
OcNOS(config)#ip vrf VRF1
OcNOS(config-vrf)#exit
OcNOS(config)#feature streaming-telemetry vrf VRF1
```

2. Create Sensor Group

Create a sensor group (`Platform`) where sensor paths will be specified for dial-out subscriptions. Specify sensor paths within the sensor group (`Platform`) to monitor the state of RAM and Hard Disk.

```
OcNOS(config)#sensor-group Platform vrf VRF1
OcNOS(telemetry-sensor-group)#sensor-path ipi:/components/component[name=RAM]/ram/state
OcNOS(telemetry-sensor-group)#sensor-path ipi:/components/component[name=HARD-DISK]/storage/state
OcNOS(telemetry-sensor-group)#exit
```

3. Create Destination Group

   Create a destination group (`Collector3`) where tunnel server settings will be configured for dial-out subscriptions. Specify the tunnel server (`gNMI Client`) IP address (`10.21.3.4`) and port (`11123`) within the destination group (`Collector3`).

```
OcNOS(config)#destination-group Collector3 vrf VRF1
OcNOS(telemetry-grpc-tunnel-group)#tunnel-server ip 10.21.3.4 port 11123
OcNOS(telemetry-grpc-tunnel-group)#exit
```

4. Create Persistent Subscription

   Create a persistent subscription (`storage`), encoding type (`JSON-IETF`), and associate it with the destination group (`Collector3`), and sensor group (`Platform`) to monitor the state of RAM and Hard Disk with a sample interval (`95 seconds`).

```
OcNOS(config)#subscription-name storage vrf VRF1
OcNOS(telemetry-subscription)#encoding json-ietf
OcNOS(telemetry-subscription)#destination-group Collector3
OcNOS(telemetry-subscription)#sensor-group Platform sample-interval 95
OcNOS(telemetry-subscription)#commit
OcNOS(telemetry-subscription)#exit
```

Streaming Telemetry Snippet Configurations on User-defined VRF

To verify the telemetry configuration and view the overall commands used for dial-out subscriptions, use the `show running-config streaming-telemetry` command on the router.

```
OcNOS#show running-config streaming-telemetry
!
feature streaming-telemetry vrf VRF1
debug telemetry gnmi enable severity debug
!
sensor-group Platform vrf VRF1
 sensor-path ipi:/components/component[name=RAM]/ram/state
 sensor-path ipi:/components/component[name=HARD-DISK]/storage/state
!
destination-group Collector3 vrf VRF1
 tunnel-server ip 10.21.3.4 port 11123
!
subscription-name storage vrf VRF1
 destination-group Collector3
 sensor-group Platform sample-interval 95
!
```

Use Case 3: Configure Telemetry on Default VRF

Note: Before configuring Dial-out, meet all Prerequisites.

1. Enable Streaming Telemetry in a default VRF on an OcNOS router.

```
OcNOS(config)#feature streaming-telemetry
```

2. Create Sensor Group

   Create a sensor group (`Platform`) where sensor paths will be specified for dial-out subscriptions. Specify sensor paths within the sensor group (`Platform`) to monitor the state of RAM and Hard Disk.

```
OcNOS(config)#sensor-group Platform
OcNOS(telemetry-sensor-group)#sensor-path ipi:/components/component[name=RAM]/ram/state
OcNOS(telemetry-sensor-group)#sensor-path ipi:/components/component[name=HARD-DISK]/storage/state
OcNOS(telemetry-sensor-group)#exit
```

3. Create Destination Group

   Create a destination group (`Collector1`) where tunnel server settings will be configured for dial-out subscriptions. Specify the tunnel server (`gNMI Client`) IP address (`10.12.101.72`) and port (`11161`) within the destination group (`Collector1`).

```
OcNOS(config)#destination-group Collector1
OcNOS(telemetry-grpc-tunnel-group)#tunnel-server ip 10.12.101.72 port 11161
OcNOS(telemetry-grpc-tunnel-group)#exit
```

4. Create Persistent Subscription

   Create a persistent subscription (`storage`), encoding type (`JSON-IETF`), and associate it with the destination group (`Collector1`), and sensor group (`Platform`) to monitor the state of RAM and Hard Disk with a sample interval (`10 seconds`).

```
OcNOS(config)#subscription-name storage
OcNOS(telemetry-subscription)#encoding json-ietf
OcNOS(telemetry-subscription)#destination-group Collector1
OcNOS(telemetry-subscription)#sensor-group Platform sample-interval 10
OcNOS(telemetry-subscription)#commit
OcNOS(telemetry-subscription)#exit
```

Streaming Telemetry Snippet Configurations on default VRF

To verify the telemetry configuration and view the overall commands used for dial-out subscriptions, use the `show running-config streaming-telemetry` command on the router.

```
OcNOS#show running-config streaming-telemetry
!
feature streaming-telemetry
debug telemetry gnmi enable severity debug
!
sensor-group Platform
 sensor-path ipi:/components/component[name=RAM]/ram/state
 sensor-path ipi:/components/component[name=HARD-DISK]/storage/state
!
destination-group Collector1
 tunnel-server ip 10.12.101.72 port 11161
```

```
!
subscription-name storage
 destination-group Collector1
 sensor-group Platform sample-interval 10
!
```

Validation

To verify persistent telemetry configurations and monitor the telemetry data transmission settings on the router, check the output of the `show streaming-telemetry persistent-subscriptions details` command.

Use Case 1: Validate Telemetry on Management VRF

```
#show streaming-telemetry persistent-subscriptions details

 Feature streaming telemetry   : Enabled

 VRF                           : management
 Platform type                 : Standard range
 Maximum sensor-paths          : 50
 Minimum sample-interval       : 90
 Number of active sensor-paths : 1 (Dial-In : 0, Dial-out : 1)
 Tunnel-server Retry-interval  : Default-60 (seconds)

 Enc-Type     : Encoding type
 SI           : Sampling Interval in seconds
 Origin:Path  : Sensor Path

 Dial-Out Subscription Details:
 ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
 1. Subscription-name    : storage2
    Status               : ACTIVE
    Enc-Type             : JSON-IETF
    Tunnel-server details:
    ~~~~~~~~~~~~~~~~~~~~~
     Destination-group         Status         Tunnel-IP:Port
     -----------------         ------         --------------
     Collector2                ACTIVE         10.21.3.4:11123
    Sensor-group details:
    ~~~~~~~~~~~~~~~~~~~~
    Sensor-group          SI            Origin:Path
    ------------         ----           -----------
        Platform            95           ipi:/components/component[name=CHASSIS]/state


 [*]-> Indicates child path learnt from parent config, not configured by user
```

Use Case 2: Validate Telemetry on User-defined VRF

```
#show streaming-telemetry persistent-subscriptions details
```

```
Feature streaming telemetry   : Enabled
 VRF                          : VRF1
 Platform type                : Standard range
 Maximum sensor-paths         : 50
 Minimum sample-interval      : 90
 Number of active sensor-paths : 2 (Dial-In : 0, Dial-out : 2)
 Tunnel-server Retry-interval  : Default-60 (seconds)

 Enc-Type     : Encoding type
 SI           : Sampling Interval in seconds
 Origin:Path  : Sensor Path

 Dial-Out Subscription Details:
 ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
 1. Subscription-name    : storage
    Status               : ACTIVE
    Enc-Type             : JSON-IETF
    Tunnel-server details:
    ~~~~~~~~~~~~~~~~~~~~~
     Destination-group         Status         Tunnel-IP:Port
     -----------------         ------         --------------
     Collector3                ACTIVE         10.21.3.4:11123
    Sensor-group details:
    ~~~~~~~~~~~~~~~~~~~~
    Sensor-group       SI       Origin:Path
    ------------       ----     -----------
        Platform       95       ipi:/components/component[name=RAM]/ram/state
                                ipi:/components/component[name=HARD-DISK]/storage/state
    [*]-> Indicates child path learnt from parent config, not configured by user
```

Use Case 3: Validate Telemetry on Default VRF

```
#show streaming-telemetry persistent-subscriptions details

 Feature streaming telemetry : Enabled

 VRF                          : default
 Platform type                : High range
 Maximum sensor-paths         : 100
 Minimum sample-interval      : 10
 Number of active sensor-paths : 2 (Dial-In : 0, Dial-out : 2)
 Tunnel-server Retry-interval  : Default-60 (seconds)

 Enc-Type     : Encoding type
 SI           : Sampling Interval in seconds
 Origin:Path  : Sensor Path

 Dial-Out Subscription Details:
 ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
 1. Subscription-name    : storage
```

```
Status                : ACTIVE
Enc-Type              : JSON-IETF
Tunnel-server details:
~~~~~~~~~~~~~~~~~~~~~
 Destination-group         Status         Tunnel-IP:Port
 -----------------         ------         --------------
 Collector1                IN-ACTIVE      10.12.101.72:11161
Sensor-group details:
~~~~~~~~~~~~~~~~~~~~
Sensor-group    SI       Origin:Path
------------   ----      -----------
    Platform     10        ipi:/components/component[name=RAM]/ram/state
                            ipi:/components/component[name=HARD-DISK]/storage/state
[*]-> Indicates child path learnt from parent config, not configured by user
```

Telemetry Subscription Invoked via gnmic Command and YAML Input

Start the gNMI collector with the `--use-tunnel-server` and `publish` options to receive the streamed gRPC responses. Execute the following command to start the gRPC tunnel server in listening mode, enabling it to accept incoming connections from gRPC tunnel clients (OcNOS target).

```
./gnmic --insecure --config <path to Tunnel-server yaml file> --use-tunnel-
server publish
```

Invoke Publish RPC on OcNOS Target

The following output represents telemetry data published by the `gnmic` command, monitoring the state of Hard Disk and RAM on the specified OcNOS router.

```
# ./gnmic --insecure --config abc.yaml --use-tunnel-server publish
2024/04/12 11:22:50.516313 [gnmic] version=dev, commit=none, date=unknown,
gitURL=, docs=https://gnmic.openconfig.net
2024/04/12 11:22:50.516377 [gnmic] using config file "abc.yaml"
2024/04/12 11:22:50.517770 [gnmic] starting output type file
2024/04/12 11:22:50.517971 [file_output:default-stdout] initialized file
output:
{"Cfg":{"FileName":"","FileType":"stdout","Format":"json","Multiline":true,"In
dent":"
","Separator":"\n","OverrideTimestamps":false,"AddTarget":"","TargetTemplate":
"","EventProcessors":null,"MsgTemplate":"","ConcurrencyLimit":1000,"EnableMetr
ics":false,"Debug":false}}
2024/04/12 11:22:50.518018 [gnmic] StartPublishCollector is invoked
2024/04/12 11:22:50.518446 [gnmic] Initializing error chan
2024/04/12 11:22:54.508410 [gnmic] tunnel server discovered target
{ID:e8:c5:7a:fe:fd:32 Type:GNMI_GNOI}
2024/04/12 11:22:54.508720 [gnmic] adding target
{"name":"e8:c5:7a:fe:fd:32","address":"e8:c5:7a:fe:fd:32","username":"root","p
assword":"****","timeout":10000000000,"insecure":true,"skip-
verify":false,"buffer-size":100,"retry-timer":10000000000,"log-tls-
secret":false,"gzip":false,"token":"","tunnel-target-type":"GNMI_GNOI"}
2024/04/12 11:22:54.508756 [gnmic] calling publishStream
2024/04/12 11:22:54.508772 [gnmic] publishStream is invoked
2024/04/12 11:22:54.508779 [gnmic] targetPublishStream is invoked
2024/04/12 11:22:54.508830 [gnmic] a.targetsChan: 0xc0004eb1a0
2024/04/12 11:22:54.508840 [gnmic] t.Config.Outputs: []
2024/04/12 11:22:54.508850 [gnmic] starting target "e8:c5:7a:fe:fd:32"
listener
```

```
2024/04/12 11:22:54.508879 [gnmic] queuing target "e8:c5:7a:fe:fd:32"
2024/04/12 11:22:54.508902 [gnmic] subscribing to target: "e8:c5:7a:fe:fd:32"
2024/04/12 11:22:54.508918 [gnmic] calling clientPublish
2024/04/12 11:22:54.508930 [gnmic] targetDialOpts: []grpc.DialOption
2024/04/12 11:22:54.508968 [gnmic] a.targetsChan: 0xc0004eb1a0
2024/04/12 11:22:54.508976 [gnmic] t.Config.Outputs: []
2024/04/12 11:22:54.509402 [gnmic] dialing tunnel connection for tunnel target
"e8:c5:7a:fe:fd:32"
Publish Request sent to e8:c5:7a:fe:fd:32{
  "source": "e8:c5:7a:fe:fd:32",
  "subscription-name": "storage",
  "timestamp": 1712920892603436151,
  "time": "2024-04-12T16:51:32.603436151+05:30",
  "updates": [
    {
      "Path": "ipi:components/component[name=HARD-DISK]/storage/state",
      "values": {
        "components/component/storage/state": {
          "free-memory": 0,
          "total-memory": 61057,
          "used-memory": 0
        }
      }
    }
  ]
}
{
  "source": "e8:c5:7a:fe:fd:32",
  "subscription-name": "storage",
  "timestamp": 1712920892603253590,
  "time": "2024-04-12T16:51:32.60325359+05:30",
  "updates": [
    {
      "Path": "ipi:components/component[name=RAM]/ram/state",
      "values": {
        "components/component/ram/state": {
          "available-high-memory": 0,
          "available-memory": 15084,
          "buffers": 101,
          "current-process-count": 227,
          "free-swap": 0,
          "shared-memory": 28,
          "total-high-memory": 0,
          "total-memory": 16010,
          "total-swap": 0,
          "used-memory": 926
        }
      }
    }
  ]
}
```

The output of the Publish RPC includes the following information:

Publish RPC Output details

Option	Description
source	Displays the MAC address associated with the management port of the target. Each gNMI device have a unique target ID, allowing the collector to distinguish responses between various targets.
subscription-name	The name of the subscription.
timestamp	The timestamp of the response.
time	The timestamp in a human-readable format.
updates	An array of updates, each containing Path and Values.
Path	The path to the published data.
values	The values of the published data.

The telemetry data output includes detailed fields for monitoring the state of the Hard Disk and RAM, offering insights into the memory and storage utilization of the OcNOS router.

1. Hard Disk State
   - **Free Memory:** The amount of free memory available on the hard disk.
   - **Total Memory:** The total capacity of memory on the hard disk.
   - **Used Memory:** The amount of memory currently in use on the hard disk.
2. RAM State
   - **Available High Memory:** The available high memory in the RAM.
   - **Available Memory:** The total available memory in the RAM.
   - **Buffers:** The number of buffer processes running in the RAM.
   - **Current Process Count:** The count of active processes in the RAM.
   - **Free Swap:** The amount of free swap space in the RAM.
   - **Shared Memory:** The shared memory usage in the RAM.
   - **Total High Memory:** The total high memory capacity in the RAM.
   - **Total Memory:** The total memory capacity in the RAM.
   - **Total Swap:** The total swap space available in the RAM.
   - **Used Memory:** The amount of memory currently in use in the RAM.

Implementation Examples

Real-time Visibility: Operators have real-time visibility into network device health and performance metrics.

Proactive Maintenance: Early detection of issues allows for proactive maintenance and troubleshooting.

Optimized Resource Allocation: Insights from telemetry data help optimize resource allocation and capacity planning.

Enhanced Network Reliability: Continuous monitoring enhances network reliability and reduces downtime.

Dial-Out Commands

The streaming telemetry dial-out mode introduces the following configuration commands.

destination-group

Use this command to create a destination-group for persistent subscriptions on the OcNOS device. The VRF parameter must match the VRF specified in the feature streaming-telemetry command. Can create and attach multiple destination-groups to activate streaming telemetry subscriptions.

Use the no form of this command to delete a destination-group.

Command Syntax

```
destination-group TUNNEL-NAME (vrf (management|NAME)|)
no destination-group TUNNEL-NAME (vrf (management|NAME)|)
```

Parameters

`TUNNEL-NAME`	Specify the name assigned to the tunnel server or collector endpoint used for telemetry data transmission.
`vrf NAME`	(Optional) Creates a destination-group for persistent subscriptions in a user-defined VRF.
`vrf management`	(Optional) Creates a destination-group for persistent subscriptions in the management VRF.

Default

None

Command Mode

Configure Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example creates a destination group named `tunnel-1` in a default VRF for transmitting telemetry data.

```
OcNOS(config)#destination-group tunnel-1
OcNOS(telemetry-grpc-tunnel-group)#commit
```

destination-group GRPC

Use this command to add a destination-group under subscriptions. Can create multiple destination-groups within a subscription mode.

Use `no` parameter of this command to remove the destination-groups.

Note: Ensure that the GRPC-GROUP-NAME is configured in the device's configuration mode before adding it to a subscription mode.

Command Syntax

```
destination-group GRPC-GROUP-NAME
no destination-group GRPC-GROUP-NAME
```

Parameters

GRPC-GROUP-NAME	Specify the name assigned to the tunnel server or collector endpoint used for telemetry data transmission.

Default

None

Command Mode

Telemetry-subscription Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

Ensure that the GRPC-GROUP-NAME (`tunnel-1`) is already configured in the current configuration mode.

```
OcNOS#configure terminal
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
 sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
destination-group tunnel-1
 tunnel-server ip 10.12.66.160 port 11163
!
subscription-name sub-1
 sensor-group stream-1 sample-interval 1000
!
!
```

The following commands illustrates how to add a destination group (`tunnel-1`) under subscription mode (`sub-1`) and verify the configuration using the show command output.

```
OcNOS(config)#subscription-name sub-1
OcNOS(telemetry-subscription)#destination-group tunnel-1
OcNOS(telemetry-subscription)#commit
OcNOS(telemetry-subscription)#exit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
 sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
```

```
destination-group tunnel-1
 tunnel-server ip 10.12.66.160 port 11163
!
subscription-name sub-1
 destination-group tunnel-1
 sensor-group stream-1 sample-interval 1000
!
!
```

encoding

Use this command to specify or modify encoding types for subscriptions in streaming telemetry.

Use `no` parameter of this command to remove the encoding option.

Note: Modifying the encoding type is not allowed for active subscriptions.

Command Syntax

```
encoding (json-ietf|json|proto)
no encoding
```

Parameters

`json-ietf`	Specifies the JSON encoding based on the IETF draft standard.
`json`	Specifies the default JSON encoding type.
`proto`	Specifies the Protocol Buffers v3 encoding type.

Default

None

Command Mode

Telemetry-subscription Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following commands demonstrate how to create a telemetry subscription named `sub-3` using the JSON encoding type.

```
OcNOS#configure terminal
OcNOS(config)#subscription-name sub-3
OcNOS(telemetry-subscription)#encoding json
OcNOS(telemetry-subscription)#commit
```

grpc-tunnel-server retry-interval

Use this command to set the interval for retry attempts when establishing a connection for the GNMI server to the tunnel-server. The VRF parameter must match the VRF specified in the feature streaming-telemetry command.

Use `no` parameter of this command to unset the retry-interval timer.

Command Syntax

```
grpc-tunnel-server retry-interval <30-3000> (vrf (management|NAME)|)
no grpc-tunnel-server retry-interval (vrf (management|NAME)|)
```

Parameters

`retry-interval <30-3000>`	Specifies the duration between retry attempts. The default retry-interval is 60 seconds.
`vrf management`	(Optional) Sets the retry-interval in the management VRF.
`vrf NAME`	(Optional) Sets the retry-interval in a user-defined VRF.

Default

None

Command Mode

Configure mode

Applicability

Introduced in the OcNOS version 6.5.2.

Example

The following configuration illustrates how to set the retry-interval timer for the gNMI server to the tunnel-server with a value of 80 seconds in a default VRF.

```
OcNOS#configure terminal
OcNOS(config)#feature streaming-telemetry
OcNOS(config)#grpc-tunnel-server retry-interval 80
OcNOS(config)#commit
```

sensor-group

Use this command to create a sensor group for persistent subscriptions in an OcNOS device. Multiple sensor groups can be created to specify the paths of interest for streaming telemetry. The VRF parameter must match the VRF specified in the feature streaming-telemetry command. These sensor groups are attached to subscriptions to activate streaming telemetry.

Use `no` parameter of this command to remove a created sensor group.

Command Syntax

```
sensor-group SENSOR-NAME (vrf (management|NAME)|)
no sensor-group SENSOR-NAME (vrf (management|NAME)|)
```

Parameters

SENSOR-NAME	Specifies the name of the sensor group.
vrf management	(Optional) Creates a sensor group in the management VRF.
vrf NAME	(Optional) Creates a sensor group in a user-defined VRF.

Default

None

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following commands demonstrate how to create a sensor group named stream-1 for persistent telemetry subscriptions in a default VRF on an OcNOS device:

```
OcNOS#configure terminal
OcNOS(config)#sensor-group stream-1
OcNOS(telemetry-sensor-group)#commit
OcNOS(telemetry-sensor-group)#exit
```

sensor-group sample-interval

Use this command to to associate a sensor group with a specific sampling interval under subscriptions for activating streaming telemetry. Multiple sensor groups can be created.

Use no parameter of this command to remove the sensor-groups from a subscription.

Note: Before adding a SENSOR-GROUP-NAME to a subscription, ensure the sensor group is already configured in the configuration mode.

Command Syntax

```
sensor-group SENSOR-GROUP-NAME sample-interval <10-3600>
no sensor-group SENSOR-GROUP-NAME
```

Parameters

SENSOR-GROUP-NAME	Specifies the name of the sensor group to be associated with the subscription.
sample-interval <10-3600>	Defines the sampling interval in seconds for the sensor group. The interval can range from 10 to 3600 seconds.

Default

None

Command Mode

Telemetry-subscription Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

Ensure that the SENSOR-GROUP-NAME (`stream-1`) is already configured in the current configuration mode.

```
OcNOS#configure terminal
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
 sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
subscription-name sub-1
!
!
```

The following commands illustrates how to add a sensor group (`stream-1`) under subscription mode (`sub-1`) and verify the configuration using the show command output.

```
OcNOS(config)#subscription-name sub-1
OcNOS(telemetry-subscription)#sensor-group stream-1 sample-interval 1000
OcNOS(telemetry-subscription)#commit
OcNOS(telemetry-subscription)#exit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
 sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
subscription-name sub-1
 sensor-group stream-1 sample-interval 1000
!
!
```

sensor-path

Use this command to add sensor paths under sensor-groups. Can add multiple sensor paths to a single sensor group.

Use `no` parameter of this command to remove sensor paths.

Command Syntax

```
sensor-path SENSOR-PATH
```

```
no sensor-path SENSOR-PATH
```

Parameters

SENSOR-PATH	Specifies the path of the telemetry data to include in the sensor group.

Default

None

Command Mode

Telemetry-sensor-group Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example demonstrates how to configure a sensor group (`stream-1`) and add multiple sensor paths to it for streaming telemetry.

```
OcNOS#configure terminal
OcNOS(config)#sensor-group stream-1
OcNOS(telemetry-sensor-group)#sensor-path ipi:/interfaces/
interface[name=eth0]/state/counters
OcNOS(telemetry-sensor-group)#sensor-path /interfaces/interface[name=xe2]/
state/counters
OcNOS(telemetry-sensor-group)#sensor-path openconfig:/interfaces/
interface[name=xe3]/state/counters
OcNOS(telemetry-sensor-group)#commit
OcNOS(telemetry-sensor-group)#exit
```

show streaming-telemetry persistent-subscriptions

Use this command to display a brief summary of the streaming-telemetry dial-out configurations. This command provides a concise view of the persistent subscription settings configured on the device.

Command Syntax

```
show streaming-telemetry persistent-subscriptions brief
show streaming-telemetry persistent-subscriptions details (SUBSCRIPTION-NAME|)
```

Parameters

SUBSCRIPTION-NAME	Displays detailed configuration information specific to the named persistent subscription.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The command output lists each persistent subscription with its associated details.

```
OcNOS#show streaming-telemetry persistent-subscriptions details

 Feature streaming telemetry : Enabled

 VRF                          : default
 Platform type                : High range
 Maximum sensor-paths         : 100
 Minimum sample-interval      : 10
 Number of active sensor-paths : 2 (Dial-In : 0, Dial-out : 2)
 Tunnel-server Retry-interval  : Default-60 (seconds)

 Enc-Type     : Encoding type
 SI           : Sampling Interval in seconds
 Origin:Path  : Sensor Path

 Dial-Out Subscription Details:
 ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
 1. Subscription-name    : State
    Status               : ACTIVE
    Enc-Type             : JSON
    Tunnel-server details:
    ~~~~~~~~~~~~~~~~~~~~~
     Destination-group         Status         Tunnel-IP:Port
     -----------------         ------         --------------
     Collector1                IN-ACTIVE      10.12.101.72:11161
    Sensor-group details:
    ~~~~~~~~~~~~~~~~~~~~
    Sensor-group      SI        Origin:Path
    ------------      ----      -----------
        storage       10        ipi:/components/component[name=RAM]/ram/state
                                ipi:/components/component[name=HARD-DISK]/storage/state
```

The following table explains the output fields.

Field	Description
Feature streaming telemetry	Marked as “Enabled” confirms that streaming telemetry is active on the device.
VRF	Specifies the VRF type.
Platform type	Displays the platform type is standard or high range.

Field	Description
Maximum sensor-paths	Shows the maximum number of sensor paths allowed. For more details, refer to Scale Scenarios section.
Minimum sample-interval	Indicates the minimum sampling interval in seconds. For more details, refer to Scale Scenarios section.
Number of active sensor-paths	Shows the total number of active sensor paths for Dial-In and Dial-Out subscriptions (Stream mode subscriptions).
Tunnel-server Default-Retry-interval	The duration between retry attempts when establishing a connection for the GNMI server to the tunnel server.
Subscription Name	Name of the persistent subscription.
Storage Status or Status	Current status of the subscription (ACTIVE or IN-ACTIVE).
Enc-Type	Encoding type used for telemetry data (JSON, JSON-IETF, Proto).
Destination Group	Define the tunnel server settings to which telemetry data is sent for dial-out subscriptions.
Sensor Group	Sensor group associated with the subscription.
Sample Interval (SI)	Sampling interval for the sensor group.
Tunnel-IP:Port	IP address and port of the tunnel server for dial-out subscriptions.
Origin:Path	The specific sensor paths that are being monitored or streamed by the telemetry system.

subscription-name

Use this command to create named subscriptions for persistent telemetry configurations in an OcNOS device. The VRF parameter must match the VRF specified in the feature streaming-telemetry command. Multiple subscriptions can be created. These subscriptions are essential for activating streaming telemetry, as they define specific settings such as associated destination groups and sensor groups.

Use `no` parameter of this command to delete a subscription.

Command Syntax

```
subscription-name NAME (vrf (management|NAME)|)
no subscription-name NAME (vrf (management|NAME)|)
```

Parameters

`subscription-name NAME`	Specifies the unique name to the persistent subscription.
`vrf NAME`	(Optional) Creates named subscriptions in a user-defined VRF.
`vrf management`	(Optional) Creates named subscriptions in the management VRF.

Default

None

Command Mode

Configure Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following command demonstrates configuring a subscription (`sub-1`) on an OcNOS device. The subscription remains `in-active` because the sensor groups and destination groups have not been added to it.

```
OcNOS#configure terminal
OcNOS(config)#subscription-name sub-1
OcNOS(telemetry-subscription)#commit
Subscription sub-1 is "in-active": sensor-group(s) and destination-group(s)
are not configured.
OcNOS(telemetry-subscription)#exit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
!
subscription-name sub-1
!
!
```

tunnel-server

Use this command to add tunnel-servers under destination groups. Can create multiple tunnel servers within a destination group.

Use `no` parameter of this command to remove a tunnel server from the destination group.

Command Syntax

```
tunnel-server ip A.B.C.D port <1-65535>
```

```
no tunnel-server ip A.B.C.D port <1-65535>
```

Parameters

`ip A.B.C.D`	Specifies the tunnel server IP address.
`port <1-65535>`	Specifies the tunnel server port-number.

Default

None

Command Mode

Telemetry-GRPC-tunnel-group Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following command demonstrates how to add a tunnel server within the destination group.

```
OcNOS#configure terminal
OcNOS(config)#destination-group tunnel-1
OcNOS(telemetry-grpc-tunnel-group)#tunnel-server ip 10.12.66.160 port 11163
OcNOS(telemetry-grpc-tunnel-group)#commit
OcNOS(telemetry-grpc-tunnel-group)#exit
```

Revised CLI Commands

The following is the revised command for telemetry.

show techsupport

- The existing syntax now includes the newly added parameter for telemetry, namely `gnmi`.
- The command `show techsupport gnmi` collects gNMI-related information for technical support. For more details, refer to the `show techsupport` command in the `Software Monitoring and Reporting` chapter in the *Sysytem Management Guide*.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Remote Procedure Call (gRPC)	gRPC protocol that uses HTTP/2 for transport and protocol buffers for serialization.
Persistent Subscription	Telemetry subscription that maintains continuous data streaming even after interruptions in connectivity.
gRPC Network Management Interface (gNMI)	A standardized protocol for network management using gRPC and protocol buffers.
Destination Group	Specifies the collector server's details and connection parameters for telemetry subscriptions.
Sensor Group	Contains sensor paths that define the specific data to be monitored and transmitted.
OpenConfig	Standardized model for network configuration and telemetry using a vendor-neutral approach.

CHAPTER 3 DHCPv6 Prefix Delegation Configuration

Overview

The prefix delegation feature facilitates the Dynamic Host Control Protocol (DHCP) server capable of assigning prefixes to DHCP clients from a global pool, enabling the Customer Premise Equipment (CPE) to learn the prefix. This feature also supports the DHCP server in assigning multiple prefixes to a single client. The user configures the IPv6 address using the learned prefix on its Local Area Network (LAN) interface with the subnet prefix. The LAN hosts are learning the subnetted prefix through Router Advertisement (RA) messages, an important Neighbor Discovery Protocol (NDP) component, enabling the device to auto-configure the number of IPv6 addresses from 1 to 64.

This feature would enable service providers to assign IP for the CPE that is acting as a router between the service providers' core network and the subscribers' internal network.

Feature Characteristics

- DHCPv6 Identity association for non-temporary addresses (IA_NA) assigns a global IPv6 address on the Wide Area Network (WAN) link. The address comes from a local pool specified in the DHCP Server.
- The Requesting Router (RR) uses the delegated prefix to define the subnet for the LAN based on the prefix received from the DHCP Server.
- The Requesting Router uses the delegated prefix to assign addresses to the LAN devices. The RR can send a Router Advertisement or the devices shall send a Router solicitation.

Benefits

The key benefits are as follows:

- This feature helps the Internet Service Providers (ISPs) to assign the dynamic IPv6 addresses to their customers automatically instead of statically assigning the address.
- This feature adds the capability to get the multiple DHCPv6 prefixes as per the customer requirement.
- This feature allows the centralized management of the IPv6 addresses.

Configuration

This section shows the configuration of the DHCPv6 prefix delegation.

Topology

The requesting router sends the prefix request to the delegating router, which sends the request to the DHCP server. The DHCP server sends the prefix to the requesting router through the delegating router. The IPv6 address is created in the requesting router by combining the prefix learned from the server and the user-defined suffix. The host receives the IPv6 address from the requesting router.

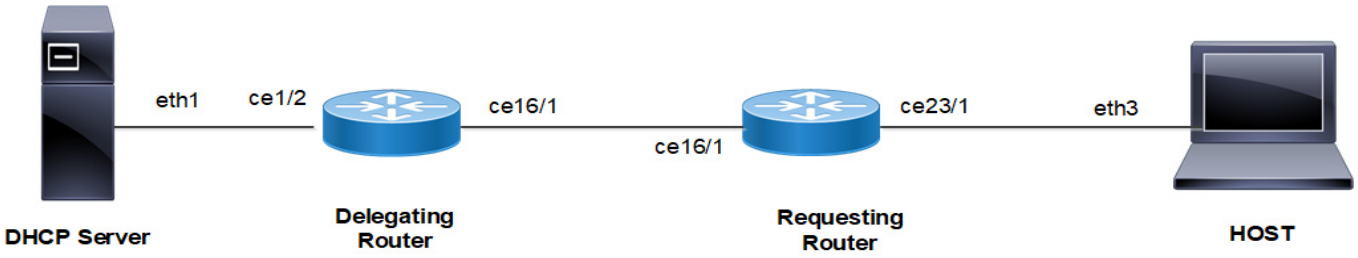

Figure 3-13: DHCPv6 Prefix Delegation Configuration

Configuring DHCP prefixes

Follow the steps to configure the DHCPv6 prefix delegation.

Configure the Delegating Router:

1. Specify the server interface address connected to the delegating router.

```
(config)#ipv6 dhcp relay address 2001:101:0:1::131
```

2. Configure the DHCPv6 up-link interface from the delegating router to the DHCPv6 server using `ipv6 dhcp relay uplink` command.

```
(config)#interface ce1/2
(config-if)#ipv6 address 2001:101:0:1::130/64
(config-if)#ipv6 dhcp relay uplink
```

3. Configure the DHCPv6 down-link interface from the delegating router to the requesting router using `ipv6 dhcp relay` command.

```
(config)#interface ce16/1
(config-if)#ipv6 address 3001:101:0:1::135/64
(config-if)#ipv6 dhcp relay
```

4. Add a static route on the delegating router to reach the host device.

```
(config)#ipv6 route ::/0 3001:101:0:1::
```

Configure the Requesting Router device:

1. In the WAN interface, configure the address prefix length option (`64`). Get the IPv6 address from the server using `ipv6 address dhcp` command. Enable the requesting router to request the prefix by using `ipv6 dhcp prefix-delegation` and configure the number of prefixes using `ipv6 dhcp client max-delegated-prefixes.`

   Note: The default value of simultaneous prefixes delegated to a single client is 8.The minimum of simultaneous prefixes delegated to a single client is 1 and the maximum is 64.

   Note: If the configured `max-delegated-prefix count` is greater than 30, then configure the lease times greater than 180 seconds.

```
(config)#interface ce16/1
(config-if)#ipv6 dhcp address-prefix-len 64
(config-if)#ipv6 address dhcp
(config-if)#ipv6 dhcp prefix-delegation PREFIX_FROM_SERVER
(config-if)#ipv6 dhcp client max-delegated-prefixes 10
```

2. In the LAN interface, configure the command `ipv6 address` to create the IPv6 address by using the DHCP prefix learned from the server and user defined suffix.

```
(config)#interface ce23/1
(config-if)#ipv6 address PREFIX_FROM_SERVER ::1:0:0:0:1/64
```

3. Add a static route on the requesting router to reach the host device.

```
(config)#ipv6 route 2001:101:0:1::/64 3001:101:0:1::135
```

Configure the HOST:

1. In the LAN interface, configure the auto-configuration to get the dynamic IPv6 address from the server.

```
(config)#interface eth3
(config-if)#ipv6 address autoconfig max-address 10
(config if)#exit
(config)#commit
```

2. Add a static route on the host to reach the server.

```
(config)#ipv6 route 2001:101:0:1::/64 3001:101:0:1::135
```

Running configurations

The running configuration for the Delegating Router is as follows:

```
#show running-config
!
ipv6 dhcp relay address 2001:101:0:1::131
!
interface ce1/2
 ipv6 address 2001:101:0:1::130/64
 ipv6 dhcp relay uplink
!
interface ce16/1
 ipv6 address 3001:101:0:1::135/64
 ipv6 dhcp relay
 commit
 end
!
```

The running configuration for the Requesting Router is as follows:

```
#show running-config
!
interface ce16/1
 ipv6 dhcp client max-delegated-prefixes 10
 ipv6 address dhcp
 ipv6 dhcp address-prefix-len 64
 ipv6 dhcp prefix-delegation PREFIX_FROM_SERVER
!
interface ce23/1
 ipv6 address PREFIX_FROM_SERVER ::1:0:0:0:1/64
 commit
 end
!
```

The running configuration for the HOST is as follows:

```
#show running-config
!
interface eth3
```

```
     ipv6 address autoconfig max-address 10
     commit
     end
    !
```

Validation

Validate the show output after configuration as shown below.

Delegating Router:

```
#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 00:03:20
C       2001:101:0:1::/64 via ::, ce1/2, 00:02:58
D       2001:db9:c0f::/48 [80/0] via fe80::eac5:7aff:fe51:723b, ce16/1, 00:00:44
C       3001:101:0:1::/64 via ::, ce16/1, 00:00:50
C       fe80::/64 via ::, ce16/1, 00:00:50
#show ipv6 dhcp pd-route
VRF : default
  2001:db9:c0a::/48 via 2001:db9:c0b::, ce16/1, (2024-03-07 06:20:43 - 2024-03-07
06:22:13)
  2001:db9:c0b::/48 via 2001:db9:c09::, ce16/1, (2024-03-07 06:20:42 - 2024-03-07
06:22:12)
  2001:db9:c0c::/48 via 2001:db9:c0d::, ce16/1, (2024-03-07 06:20:39 - 2024-03-07
06:22:09)
  2001:db9:c0d::/48 via 2001:db9:c0e::, ce16/1, (2024-03-07 06:20:38 - 2024-03-07
06:22:08)
  2001:db9:c0e::/48 via 2001:db9:c0f::, ce16/1, (2024-03-07 06:20:37 - 2024-03-07
06:22:07)
  2001:db9:c0f::/48 via fe80::eac5:7aff:fe51:723b, ce16/1, (2024-03-07 06:20:36 - 2024-
03-07 06:22:06)
  2001:db9:c05::/48 via 2001:db9:c06::, ce16/1, (2024-03-07 06:20:45 - 2024-03-07
06:22:15)
  2001:db9:c06::/48 via 2001:db9:c0a::, ce16/1, (2024-03-07 06:20:44 - 2024-03-07
06:22:14)
  2001:db9:c08::/48 via 2001:db9:c0c::, ce16/1, (2024-03-07 06:20:40 - 2024-03-07
06:22:10)
  2001:db9:c09::/48 via 2001:db9:c08::, ce16/1, (2024-03-07 06:20:41 - 2024-03-07
06:22:11)
#show ipv6 dhcp relay
IPv6 DHCP relay service is Enabled.
VRF Name: default
  DHCPv6 Servers configured:
  2001:101:0:1::131
```

```
  DHCPv6 IA_PD Route injection: Enabled
  DHCPv6 Duplicate Clients detection: Disabled
  Interface                Uplink/Downlink
  ---------                -------------
  ce16/1                   Downlink
  ce1/2                    Uplink
```

Requesting Router:

```
#show ipv6 dhcp interface

ce16/1 is in client mode
        prefix name: PREFIX_FROM_SERVER
        learned prefix: 2001:db9:c05::/48
        preferred lifetime 0, valid lifetime 60
        interfaces using the learned prefix
                ce23/1    2001:db9:c0f:1::1
                ce23/1    2001:db9:c0e:1::1
                ce23/1    2001:db9:c0d:1::1
                ce23/1    2001:db9:c0c:1::1
                ce23/1    2001:db9:c08:1::1
                ce23/1    2001:db9:c09:1::1
                ce23/1    2001:db9:c0b:1::1
                ce23/1    2001:db9:c0a:1::1
                ce23/1    2001:db9:c06:1::1
                ce23/1    2001:db9:c05:1::1

#show interface ce23/1
Interface ce23/1
  Flexport: Non Control Port (Active)
  Hardware is ETH  Current HW addr: e8c5.7a51.722e
  Physical:e8c5.7a51.722e  Logical:(not set)
  Forward Error Correction (FEC) configured is Auto (default)
  FEC status is N/A
  Port Mode is Router
  Protected Mode is Promiscuous
  Interface index: 10017
  Metric 1 mtu 1500 duplex-full  link-speed 10g
  Debounce timer: disable
  ARP ageing timeout 1500
  <UP,BROADCAST,RUNNING,ALLMULTI,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  Bandwidth 10g
  Maximum reservable bandwidth 10g
    Available b/w at priority 0 is 10g
    Available b/w at priority 1 is 10g
    Available b/w at priority 2 is 10g
    Available b/w at priority 3 is 10g
```

```
    Available b/w at priority 4 is 10g
    Available b/w at priority 5 is 10g
    Available b/w at priority 6 is 10g
    Available b/w at priority 7 is 10g
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: Never
  inet6 2001:db9:c05:1::1/64
  inet6 2001:db9:c06:1::1/64
  inet6 2001:db9:c08:1::1/64
  inet6 2001:db9:c09:1::1/64
  inet6 2001:db9:c0a:1::1/64
  inet6 2001:db9:c0b:1::1/64
  inet6 2001:db9:c0c:1::1/64
  inet6 2001:db9:c0d:1::1/64
  inet6 2001:db9:c0e:1::1/64
  inet6 2001:db9:c0f:1::1/64
  inet6 fe80::eac5:7aff:fe51:722e/64
  ND router advertisements are sent approximately every 561 seconds
  ND next router advertisement due in 517 seconds.
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
  5 minute input rate 82 bits/sec, 0 packets/sec
  5 minute output rate 191 bits/sec, 0 packets/sec
  RX
    unicast packets 0 multicast packets 25 broadcast packets 0
    input packets 25 bytes 2862
    jumbo packets 0
    undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
    input error 0
    input with dribble 0 input discard 0
    Rx pause 0
  TX
    unicast packets 0 multicast packets 38 broadcast packets 0
    output packets 38 bytes 5540
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0
```

HOST:

```
#show ipv6 interface eth3 brief
Interface            IPv6-Address                                          Admin-Status
eth3                 2001:db9:c05:1:923c:b3ff:fe90:9fa9
                     2001:db9:c06:1:923c:b3ff:fe90:9fa9
                     2001:db9:c08:1:923c:b3ff:fe90:9fa9
                     2001:db9:c09:1:923c:b3ff:fe90:9fa9
                     2001:db9:c0a:1:923c:b3ff:fe90:9fa9
                     2001:db9:c0b:1:923c:b3ff:fe90:9fa9
                     2001:db9:c0c:1:923c:b3ff:fe90:9fa9
```

```
                    2001:db9:c0d:1:923c:b3ff:fe90:9fa9
                    2001:db9:c0e:1:923c:b3ff:fe90:9fa9
                    2001:db9:c0f:1:923c:b3ff:fe90:9fa9
                    fe80::923c:b3ff:fe90:9fa9                          [up/up]
```

DHCP Multiple Prefix Delegation Command

The DHCPv6 Prefix Delegation introduces the following configuration command.

ipv6 dhcp client max-delegated-prefixes

Use this command to configure multiple DHCPv6 prefix delegation for a single client.

Command Syntax

```
ipv6 dhcp client max-delegated-prefixes <1-64>
```

Parameters

`max-delegated-prefixes <1-64>`	Specifies the number of prefixes need for a DHCP client. Default number of DHCP prefixes are 8.

Default

None

Command Mode

Interface mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

This example shows how to configure multiple DHCPv6 prefix delegation for a single client:

```
RR#configure terminal
RR#(config)#interface ce16/1
RR#(config-if)#ipv6 dhcp address-prefix-len 64
RR#(config-if)#ipv6 address dhcp
RR#(config-if)#ipv6 dhcp prefix-delegation PREFIX_FROM_SERVER
RR#(config-if)#ipv6 dhcp client max-delegated-prefixes 10
RR#(config-if)#exit
RR#(config)#commit
```

Revised CLI Commands

The following command is revised:

ipv6 address autoconfig

The existing syntax now includes the newly added parameter `(max-address <1-64>|)`. For more details, refer to ipv6 dhcp prefix-delegation command in the DHCPv6 Prefix Delegation Commands chapter in the *System Management Guide*.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Border Network Gateway (BNG)	Border Network Gateway is a critical component in the telecommunication network that serves as the entry and exit point between the ISP and the global network.
Customer Premises Equipment (CPE)	Customer Premises Equipment is a networking device located on the customer premises. It is present on the edge of the service provider network, which connects the customer devices to the service provider network.
Delegating Router (DR)	Delegating Router is a network device that delegates the IPv6 address prefixes to the downstream devices.
Identity association for non-temporary addresses (IA_NA)	Identity association for non-temporary addresses is a unique identifier associated with a set of IPv6 addresses assigned to client devices permanently or for a long time.
Local Area Network (LAN)	Local Area Network is a network of devices in a small area that may include a building or home.
Neighbor Discovery Protocol (NDP)	Neighbor Discovery Protocol is a crucial protocol in the IPv6 networks, helping establish the communication and auto-configuration to run the devices in the local network segment seamlessly.
Neighbor Discovery Router Advertisement (NDRA)	Neighbor Discovery Router Advertisement facilitates a network device to advertise the routing information with the neighboring device so that the neighboring devices take the forwarding decision in dynamic routing.
Router Advertisement (RA)	Router Advertisement is a critical component in the IPv6 network. The router sends a message to the devices connected to the LAN to communicate its presence and share the configurations with the LAN host.
Requesting Router (RR)	Requesting Router is a network device that requests the IPv6 address prefixes to the DHCP server to share it with the downstream devices.
Router Solicitation (RS)	Router Solicitation is a component of the neighbor discovery protocol in the IPv6 network where the host sends a message to discover routers in the local area. When a router receives RS, it responds to the host with RA, which includes the configuration.
Wide Area Network (WAN)	Wide Area Network refers to large network that includes multiple LANs and spans over a large geographical area.

CHAPTER 4 Configure SRv6 with EVPN ELAN

Overview

The Ethernet Virtual Private Network - Ethernet LAN (EVPN ELAN) SRv6 feature integrates Segment Routing over IPv6 (SRv6) technology with EVPN signaling mechanisms to deliver multipoint-to-multipoint VPN services efficiently. To overcome the limitations of traditional L2VPN technologies such as Virtual Private LAN Services (VPLS), SRv6 EVPN ELAN utilizes BGP extensions and integrates the control planes for multiple VPN services. This approach separates forwarding and control planes, enabling a more efficient and effective network architecture.

Feature Characteristics

- Utilizes BGP extensions for MAC address learning and advertisement, enhancing control-plane based MAC learning.
- Supports local MAC address learning using ARP and remote MAC/IP address learning through MAC/IP advertisement routes.
- Advertises MAC/IP routes to reduce broadcast traffic volume and save bandwidth resources.
- Supports Inclusive Multicast Ethernet Tag Route (IMET) routes for efficient delivery of Broadcast, unknown Unicast, and Multicast (BUM) traffic.

Benefits

- Enhances network scalability and efficiency by moving MAC address learning to the control plane.
- Reduces network complexity and signaling messages by leveraging BGP for PE communication.
- Optimizes resource consumption by locally storing MAC and IP address information.
- Enables fast convergence and traffic balancing, improving overall network performance.

Prerequisites

Compatible network devices supporting SRv6 and EVPN technologies.

Configuration

Configure EVPN ELAN services with the SRv6 transport option, enabling enhanced scalability, flexibility, and operational efficiency.

The following configuration enables EVPN ELAN service specific to SRv6 transport.

Topology

The topology includes with edge and intermediate nodes, utilizing SRv6 functionality, and various routing protocols to ensure efficient communication and service delivery within the provider network.

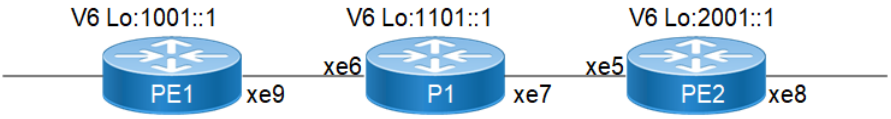

Figure 4-14: SRV6 EVPN ELAN Topology

Provider Edge Nodes (PE1 and PE2):

These intermediate nodes within the provider network may or may not be SRv6-capable routers.

Perform the following steps to configure SRv6 EVPN functionality on PE nodes with ISIS as IGP, appropriate MAC-VRF, BGP and EVPN EVI settings:

1. Configure Loopback Interfaces:
   - Access interface configuration mode for the loopback interface(`interface lo`).
   - Assign an IPv6 address to the loopback interface using the ipv6 address command followed by the desired IPv6 address and subnet mask (`ipv6 address 1001::1/128`).
   - Configure OSPF for IPv6 on the loopback interface using the ipv6 router ospf command, specifying the OSPF area, tag, and instance ID (`ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0`).
   - Configure IS-IS for IPv6 on the loopback interface using the ipv6 router isis command, specifying the IS-IS process ID (`ipv6 router isis 1`).

     ```
     PE1(config)#interface lo
     PE1(config-if)#ipv6 address 1001::1/128
     PE1(config-if)#ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
     PE1(config-if)#ipv6 router isis 1
     PE1(config-if)#exit
     ```

2. Configure Network interfaces:
   - Access interface configuration mode for the desired network interface `(interface xe9)`.
   - Assign an IPv6 address to the interface using the ipv6 address command followed by the desired IPv6 address and subnet mask `(ipv6 address cafe:1:1::1/64)`.
   - Configure the MTU for the interface `(mtu 9216)`.
   - Configure OSPF for IPv6 on the interface using the ipv6 router ospf command, specifying the OSPF area, tag, and instance ID `(ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0)`.
   - Configure IS-IS for IPv6 on the interface using the ipv6 router isis command, specifying the IS-IS process ID `(ipv6 router isis 1)`.

     ```
     PE1(config)#interface xe9
     PE1(config-if)#ipv6 address cafe:1:1::1/64
     PE1(config-if)#mtu 9216
     PE1(config-if)#ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
     PE1(config-if)#ipv6 router isis 1
     ```

3. In Global configuration mode, perform the following:
   - Enable EVPN SRv6 for EVPN on the router, allowing for flexible and scalable IPv6-based service delivery.

     ```
     PE1(config)# evpn srv6 enable
     ```

   - Configure global IPv6 address for SRv6 functionality in the EVPN on the router:

     ```
     PE1(config)# evpn srv6 ip-global 1001::1
     ```

   - Configure QOS.

     ```
     PE1(config)#qos enable
     ```

- Define SRv6 locators to be used in the EVPN configuration.

```
PE1(config)# segment-routing srv6
PE1(config-srv6)# locators
PE1(config-locator)# locator PE1_locator
PE1(config-locator)# prefix 1001::/64
PE1(config-locator)# exit-locators
PE1(config-srv6)# exit-srv6
```

4. Configure ISIS Settings:
   - Access ISIS configuration mode and provide the ISIS process ID (`router isis 1`).
   - Specify the ISIS routing level using the is-type (`is-type level-2-only`).
   - Configure the metric-style wide (`metric-style wide`)
   - Enable dynamic hostname assignment.
   - Configure the NET address (`net 49.0001.0000.0000.0001.00`).
   - Enter address-family configuration mode for IPv6 (`address-family ipv6`).
   - Configure segment routing with SRv6 (`segment-routing srv6`)

```
PE1(config)#router isis 1
PE1(config-router)#is-type level-2-only
PE1(config-router)#metric-style wide
PE1(config-router)#dynamic-hostname
PE1(config-router)#net 49.0001.0000.0000.0001.00
PE1(config-router#address-family ipv6
PE1(config-router-af)#segment-routing srv6
PE1(config-router-af-srv6)#srv6-locator PE1_locator
PE1(config-router-af-srv6)#exit-srv6
PE1(config-router-af)# exit-address-family
```

5. Perform the BGP Configuration:

```
PE1(config)#router bgp 65010
PE1(config-router)#bgp router-id 1.1.1.1
PE1(config-router)#neighbor 2001::1 remote-as 65010
PE1(config-router)#neighbor 2001::1 update-source lo
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor 2001::1 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#exit
```

6. Create MAC VRF:

```
PE1(config)#mac vrf PE1_PE2_ELAN
PE1(config-vrf)#rd 1.1.1.1:2000
PE1(config-vrf)#route-target both 2000:2000
```

7. Define the EVI instance and SRv6 for the EVI with the MAC VRF Mapping specified locator:

```
PE1(config)#evpn srv6 id 2000
PE1(config)#host-reachability-protocol evpn-bgp PE1_PE2_ELAN
PE1(config)# locator PE1_locator
PE1(config)# exit
PE1(config)#interface xe6.2000 switchport
```

```
  PE1(config-if)#encapsulation dot1q 2000
  PE1(config-if)#mtu 9216
  PE1(config-if)#access-if-evpn
  PE1(config-access-if)#map vpn-id 2000
```

Configuration Snapshot: SRv6 EVPN Single-Homing on PE1

```
evpn srv6 enable
!
mac vrf PE1_PE2_ELAN
 rd 1.1.1.1:2000
 route-target both 2000:2000
!
qos enable
!
evpn srv6 ip-global 1001::1
!
evpn srv6 id 2000
 host-reachability-protocol evpn-bgp PE1_PE2_ELAN
 locator PE1_locator
!
hostname PE1
!
router-id 1.1.1.1
!
segment-routing
 srv6
  locators
   locator PE1_locator
    prefix 1001::/64
   exit-locator
   !
  exit-locators
 !
 exit-srv6
 !
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
 ipv6 address 1001::1/128
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
interface xe6
 mtu 9216
!
interface xe6.2000 switchport
 encapsulation dot1q 2000
 mtu 9216
 access-if-evpn
  map vpn-id 2000
!
interface xe9
 ipv6 address cafe:1:1::1/64
```

```
 mtu 9216
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
router isis 1
 is-type level-2-only
 metric-style wide
 dynamic-hostname
 net 49.0001.0000.0000.0001.00
 !
 address-family ipv6
 segment-routing srv6
  srv6-locator PE1_locator
 exit-srv6
 !
 exit-address-family
!
router bgp 65010
 bgp router-id 1.1.1.1
 neighbor 2001::1 remote-as 65010
 neighbor 2001::1 update-source lo
 !
 address-family l2vpn evpn
 neighbor 2001::1 activate
 exit-address-family
 !
 exit
!
```

Configuration Snapshot: SRv6 EVPN ELAN Single-Homing on P1

```
hostname P1
!
qos enable
!
router-id 1.1.1.11
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
 ipv6 address 1101::1/128
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
interface xe6
 ipv6 address cafe:1:11::2/64
 mtu 9216
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
interface xe7
 ipv6 address cafe:11:21::1/64
 mtu 9216
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
```

```
router isis 1
 is-type level-2-only
 metric-style wide
 dynamic-hostname
 net 49.0001.0000.0000.0011.00
 !
 address-family ipv6
 exit-address-family
!
```

Configuration Snapshot: SRv6 EVPN ELAN Single-Homing on PE2

```
evpn srv6 enable
!
mac vrf PE1_PE2_ELAN
 rd 1.1.1.2:2000
 route-target both 2000:2000
!
qos enable
!
evpn srv6 ip-global 2001::1
!
evpn srv6 id 2000
 host-reachability-protocol evpn-bgp PE1_PE2_ELAN
 locator PE2_locator
!
hostname PE2
!
router-id 1.1.1.2
!
segment-routing
 srv6
  locators
   locator PE2_locator
    prefix 2001::/64
   exit-locator
   !
  exit-locators
 !
 exit-srv6
 !
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
 ipv6 address 2001::1/128
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
interface xe5
 ipv6 address cafe:11:21::2/64
 mtu 9216
 ipv6 router ospf area 0.0.0.0 tag 1 instance-id 0
 ipv6 router isis 1
!
```

```
interface xe8
 mtu 9216
!
interface xe8.2000 switchport
 encapsulation dot1q 2000
 access-if-evpn
  map vpn-id 2000
!
router isis 1
 is-type level-2-only
 metric-style wide
 dynamic-hostname
 net 49.0001.0000.0000.0002.00
 !
 address-family ipv6
 segment-routing srv6
  srv6-locator PE2_locator
 exit-srv6
 !
 exit-address-family
!
router bgp 65010
 bgp router-id 1.1.1.2
 neighbor 1001::1 remote-as 65010
 neighbor 1001::1 update-source lo
 !
 address-family l2vpn evpn
 neighbor 1001::1 activate
 exit-address-family
 !
 exit
!
```

Validation

PE1

- The following show outputs displays the ISISv6 neighbour and routing information of the PE1.

```
PPE1#sh clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
P1             xe9         80a2.355b.7008      Up     24        L2   IS-IS
PE1#

PE1#sh clns neighbors detail

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
```

```
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol
P1             xe9         80a2.355b.7008      Up     21        L2   IS-IS
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 00:53:18
  Area Address(es): 49.0001
  IPv6 Address(es): fe80::82a2:35ff:fe5b:7008
  Level-2 Protocols Supported: IPv6
  Adjacency advertisement: Advertise

PE1#sh ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 00:56:00
C       1001::1/128 via ::, lo, 00:55:11
C       1001::6001:0:0:0/128, SRV6 END.X SID
        via fe80::82a2:35ff:fe5b:7008, xe9, 00:53:22
i L2    1101::1/128 [115/20] via fe80::82a2:35ff:fe5b:7008, xe9, 00:53:07
i L2    2001::/64 [115/21] via fe80::82a2:35ff:fe5b:7008, xe9, 00:37:00
i L2    2001::1/128 [115/30] via fe80::82a2:35ff:fe5b:7008, xe9, 00:37:00
C       cafe:1:1::/64 via ::, xe9, 00:53:23
i L2    cafe:1:11::/64 [115/20] via fe80::82a2:35ff:fe5b:7008, xe9, 00:53:07
i L2    cafe:2:3::/64 [115/30] via fe80::82a2:35ff:fe5b:7008, xe9, 00:37:00
i L2    cafe:11:3::/64 [115/20] via fe80::82a2:35ff:fe5b:7008, xe9, 00:53:07
i L2    cafe:11:21::/64 [115/20] via fe80::82a2:35ff:fe5b:7008, xe9, 00:38:40
C       fe80::/64 via ::, xe9, 00:53:23
PE1#
             2001::1                   None
```

- The following show outputs displays the BGP validation for EVPN ELAN.

```
PE1#sh bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 65010
BGP table version is 27
1 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down
State/PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
2001::1                4 65010       151       185     27     0     0  00:24:07
1      0      0      1      0         0

Total number of neighbors 1
```

```
Total number of Established sessions 1
PE1#sh ip bgp neighbors
BGP neighbor is 2001::1, remote AS 65010, local AS 65010, internal link, peer index: 7
  BGP version 4, local router ID 1.1.1.1, remote router ID 1.1.1.2
  BGP state = Established, up for 00:24:12
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family L2VPN EVPN: advertised and received
  Received 148 messages, 4 notifications, 0 in queue
  Sent 179 messages, 6 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

For address family: L2VPN EVPN  BGP table version 27, neighbor version 27
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  Accepted AD:0 MACIP:0 MCAST:1 ESI:0 PREFIX:0
  3 announced prefixes

 Connections established 9; dropped 8
Local host: 1001::1, Local port: 179
Foreign host: 2001::1, Foreign port: 45691
TCP MSS: (0), Advertise TCP MSS: (9156), Send TCP MSS: (9156),  Receive TCP MSS: (536)
Sock FD : (28)
Nexthop: 1.1.1.1
Nexthop global: 1001::1
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:24:12, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

- The following show outputs displays the SRv6 EVPN ELAN validation.

```
 PE1#show segment-routing srv6 services
Status codes: > - installed, * - selected, T - Uses service-mapped tunnel
L3VPN:

EVPN:
Service Flags vrf        local-evpn-id   remote-evpn-id     SID
Nexthop                   SRv6-Policy-Name
ELAN    >     PE1_PE2_ELAN 2000           NA                 2001::4:0:0:0
2001::1                   None
PE1#show segment-routing srv6 services evpn
Status codes: > - installed, * - selected, T - Uses service-mapped tunnel
Service Flags vrf        local-evpn-id   remote-evpn-id     SID
Nexthop                   SRv6-Policy-Name
ELAN    >     PE1_PE2_ELAN 2000           NA                 2001::4:0:0:0
2001::1                   None
```

```
PE1#show segment-routing srv6 sid
SRv6 Segment ID table:
SID                     Operation    Nexthop                    Originator
+---------------------+------------+--------------------------+----------
1001::3:0:0:0           END.DT2U     ::                         evpn:2000
1001::4:0:0:0           END.DT2M     ::                         evpn:2000
1001::801:0:0:0         END[usd]     ::                         nsm
1001::1001:0:0:0        END[usp]     ::                         nsm
1001::2001:0:0:0        END[psp]     ::                         nsm
1001::6001:0:0:0        END.X[psp]   fe80::82a2:35ff:fe5b:7008isis
PE1#
PE1#show hsl srv6 evpn
TABLE: SRV6 EVPN Table
+--------+---------------------+--------------+--------+--------------------------+----
---------------------+------+----------+--------------------------+----------+
|  EVPN  |     DESTINATION     |  POLICY-ID/  |  OUT   |         NEXTHOP          |
EVPN UC SID        |CACHE |    UC    |       EVPN MC SID        |   VSI    |
|   ID   |       FEC       | TYPE/NHLFE-ID | IFNAME |                          |
|STATUS|  FEC_ID  |                           |         |
+--------+---------------------+--------------+--------+--------------------------+----
---------------------+------+----------+--------------------------+----------+
 2000   2001::1        0  /PRI/4   xe9   fe80::82a2:35ff:fe5b:7008 ::
2001::4:0:0:0              4154
PE1#
PE1#show evpn srv6 id 2000
EVPN-SRv6 Information
=================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VPN-ID   EVI-Name     EVI-Type Type Interface ESI                              VLAN      DF-
Status Src-Addr         Dst-Addr
_____________________________________________________________________________________
_____________________________________
2000     ----         L2       NW   ----      ----                             ----      -
---      1001::1          2001::1
2000     ----         --       AC   xe6.2000  --- Single Homed Port ---     ----      -
---      ----             ----

Total number of entries are 2

Note: Refer sub-interface config for VLAN information.
PE1#show evpn srv6  tunnel summary

Total number of entries: 1 [Installed: 1, Resolved: 0, Unresolved: 0]
PE1#show evpn srv6  tunnel sid
EVPN-SRV6 Network tunnel SID's
 Evpn service type: ELAN, evi: 2000, evi-name: , status: Installed
   PE IP: 2001::1
    Tunnel information
     local UC-SID: 1001::3:0:0:0, local MC-SID: 1001::4:0:0:0
```

```
     remote UC-SID: ::, remote MC-SID: 2001::4:0:0:0
    Tunnel policy mapped: --

Total number of entries are 1
PE1#
```

Implementation Examples

The SRv6 technology can be used to implement different use cases, such as MAC/IP Advertisement Route and IMET Route over SRv6 Core Propagation. In both cases, the SRv6-enabled routers learn MAC address information from the packets they receive and cache it in the forwarding tables, which helps optimize resource consumption and improve overall network performance. The SRv6 technology also helps reduce network complexity by leveraging BGP for PE communication and enables fast convergence and traffic balancing.

CLI Commands

The `EVPN ElAN SRv6` introduces the following configuration commands:

- `evi-name`
- `evpn srv6 mac-ageing-time`
- `arp-nd refresh timer`
- `mac-holdtime`
- `show evpn srv6`
- `show evpn srv6 arp-cache`
- `show evpn srv6 mac-table`
- `show evpn srv6 nd-cache`
- `show evpn srv6 route-count`
- `show evpn srv6 static host state`

evi-name

Use this command to name the EVPN MPLS ID.

Use `no` parameter of this command to remove the name of the EVPN SRv6 ID.

Command Syntax

```
evi-name <WORD>
no evi-name
```

Parameters

`WORD`	EVI name of max size 10 characters and should not be only numeric.

Default

None

Command Mode

EVPN SRv6 mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to enable srv6 for EVPN.

```
#configure terminal
(config)#evpn srv6 id 3
(config-evpn-srv6)#evi-name ELAN
(config-evpn-srv6)#exit
```

evpn srv6 mac-ageing-time

Use this command to set the dynamically learned MAC aging time.

Use `no` parameter of this command to set the age out the MACs in hardware to its default.

Command Syntax

```
evpn srv6 mac-ageing-time <10-572>
no evpn srv6 mac-ageing-time
```

Parameters

`mac-ageing-time<10-572>`	EVI name of max size 10 characters and should not be only numeric.

Default

Age out time to 300 seconds

Command Mode

Config mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to configure `evpn srv6 mac-ageing-time:`

```
#configure terminal
(config)#evpn srv6 mac-ageing-time 10
```

arp-nd refresh timer

Use this command to configure aging out the arp-cache and nd-cache entries for given time multiplied by 3 in seconds.

Use `no` parameter of this command to remove the configuration.

Note:

- Not applicable for the AC port which is mapped with ELINE/Xconnect Service.
- After this timer interval, it sends out ARP to revalidate and 3 times of this would lead to removal of the dynamic entry.

Command Syntax

```
evpn srv6 arp-nd refresh-timer <3-190>
no evpn srv6 arp-nd refresh-timer
```

Parameters

`arp-nd refresh-timer<3-190>`	Sets the refresh timer value for ARP and ND cache entries on a networking device.

Default

Disabled

Command Mode

Evpn mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to configure `evpn srv6 arp-nd refresh-timer:`

```
(config)#evpn srv6 arp-nd refresh-timer 100
(config)#no evpn srv6 arp-nd refresh-timer
```

mac-holdtime

Use this command to set the MAC hold time for a MAC/IP or MAC.

The feature holds the MAC in hardware until BGP has withdrawn from the neighbours. This helps to reduce the flooding to other access ports.This setting applies when the L2 Subifp is shut down, the physical port on which the access port is down, or the access port is removed from the VNID using the no form of the map vnid command.When the MAC hold time is configured as -1, then the MAC is not removed from the hardware and is also not withdrawn from EVPN BGP.

Use the `no` form of this command to remove the MAC hold time for the MAC/IP or MAC

Note: When a MAC address enters the discard state, traffic associated with it is dropped. This rule applies exclusively to MAC addresses or MAC-IP pairs configured manually.

Command Syntax

```
mac-holdtime <-1-300>
no mac-holdtime
```

Parameters

`<-1-300>`	MAC hold time in seconds. Specify -1 to never expire state.

Default

Zero second

Command Mode

EVPN SRv6 mode and ACC_IF mode.

Note: When set in both modes, the preference is given to the ACC_IF mode value for the corresponding access port.

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to configure `mac-holdtime` for `evpn srv6`:

```
#configure terminal
(config)#evpn srv6 id 3
(config-evpn-srv6) #mac-holdtime -1
(config-evpn-srv6) #exit
```

show evpn srv6

Use this command to display the EVPN Information.

Command Syntax

```
show evpn srv6 ((tunnel (| sid | summary) | id <1-16777215>)|)
```

Parameters

`tunnel sid`	Displays Segment Identifier (SID) used in Segment Routing (SR) networks to identify a tunnel.
`tunnel summary`	Provides a summarized view of SRv6 configurations and statuses.
`tunnel id <1-16777215>`	Displays information related to the specified SRv6 tunnel or SID identified by its numerical ID. The ID range is from 1 to 16777215.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to display the show output of `evpn srv6 tunnel`.

```
PE1# show evpn srv6 tunnel sid
EVPN-SRV6 Network tunnel SID's
 Evpn service type: ELAN, evi: 10, evi-name: , status: Installed
   PE IP: 2001::3
    Tunnel information
     local UC-SID: cafe:aaaa:1:0:2::, local MC-SID: cafe:aaaa:1:0:3::
     remote UC-SID: cafe:aaaa:3:0:2::, remote MC-SID: cafe:aaaa:3:0:3::
    Tunnel policy mapped: --
 Evpn service type: ELAN, evi: 10, evi-name: , status: Installed
   PE IP: 2001::2
    Tunnel information
     local UC-SID: cafe:aaaa:1:0:2::, local MC-SID: cafe:aaaa:1:0:3::
     remote UC-SID: cafe:aaaa:2:0:2::, remote MC-SID: cafe:aaaa:2:0:3::
    Tunnel policy mapped: --

Total number of entries are 2
```

show evpn srv6 arp-cache

Use this command to display the ARP cache information.

Command Syntax

```
show evpn srv6 arp-cache (evid <1-16777215>|) (summary |)
```

Parameters

`arp-cache`	Displays ARP cache information for all EVPN instances.
`evid <1-16777215>`	Displays ARP cache information specific to the EVPN instance identified by its Ethernet Segment Identifier (EVID). The EVID range is from 1 to 16777215.
`summary`	Provides a summarized view of the ARP cache information.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to display the show output of `evpn srv6 arp-cache`

```
PE1#show evpn srv6 arp-cache
SRV6-EVPN ARP-CACHE Information
================================
EVPN-ID   Ip-Addr          Mac-Addr        Type        Age-Out   Retries-Left
_____________________________________________________________________________
10        7.7.7.7          0020.9400.0004 Static Local   ----
10        192.85.1.3       0010.9400.0003 Dynamic Remote ----
10        192.85.1.4       0010.9400.0004 Dynamic Local  ----
Total number of entries are 3
```

show evpn srv6 mac-table

Use this command to display the host MAC address table.

Command Syntax

```
show evpn srv6 mac-table (hardware |) (evid <1-16777215>|) (summary |)
```

Parameters

`mac-table`	Displays the EVPN SRv6 MAC address table.
`evid <1-16777215>`	Specifies the EVPN Instance Identifier (EVI) for which you want to display the SRv6 MAC table information. The range for the EVI ID is from 1 to 16777215.
`hardware`	Displays Host mac addresses table from hardware.
`summary`	Provides a summarized view of Host mac addresses table from hardware.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to display the show output of `evpn srv6 mac-table`

```
PE1#show evpn srv6 mac-table
================================================================================
================================================================================
                                                        EVPN SRV6 MAC Entries
================================================================================
================================================================================
VNID    Interface VlanId  In-VlanId Mac-Addr    VTEP-Ip/ESI
Type            Status     MAC move AccessPortDesc
________________________________________________________________________________
________________________________________________________________________________

10      ----    ----    ----    0001.9400.0003 2001::3
Static Remote   -------    0        -------
```

```
10      ----    ----    ----    0011.9400.0003 2001::3
Static Remote   -------    0        -------
10      ----    ----    ----    0011.9401.0003 2001::3
Static Remote   -------    0        -------
10      xe29.100 ----    ----    0020.9400.0003 2001::1
Static Local    -------    0        -------
10      xe29.100 ----    ----    0030.9400.0003 2001::1
Static Local    -------    0        -------

Total number of entries are : 5
```

show evpn srv6 nd-cache

Use this command to display the Neighbor Discovery (ND) cache information.

Command Syntax

```
show evpn srv6 nd-cache (evid <1-16777215>|) (summary |)
```

Parameters

`nd-cahce`	Displays the EVPN SRv6 ND table.
`evid<1-16777215>`	Displays ND cache information specific to the EVPN instance identified by its Ethernet Segment Identifier (EVID). The EVID range is from 1 to 16777215.
`Summary`	Provides a summarized view of the ND cache information.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to display the show output of `evpn srv6 nd-cache:`

```
PE1#show evpn srv6 nd-cache
SRV6-EVPN ND-CACHE Information
==============================
EVPN-ID  Ip-Addr                        Mac-Addr          Type         Age-Out
Retries-Left

_____________________________________________________________________________
_________________
10       1111::33                       0011.9401.0003 Static Remote    ----
10       2222::22                       0011.9401.0002 Static Remote    ----
Total number of entries are 2
```

show evpn srv6 route-count

Use this command to display the EVPN active route (MAC-IP.MAC-IPv6 and MAC-only) count information.

Command Syntax

```
show evpn srv6 route-count (|evid <1-16777215>)
```

Parameters

`evid <1-16777215>`	Displays the count of SRv6 routes specific to the EVPN instance identified by its EVID. The EVID range is from 1 to 16777215.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to display the show output of `evpn srv6 route-count`

```
PE1#show evpn srv6 route-count
EVPN-SRv6 Active route count information
==================================
Max supported route count   : 131072
Active route count: 8

-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 10       8        4        2        2

Total number of entries are 1
PE1#
PE1#show evpn srv6 route-count evid 10
EVPN-SRv6 Active route count information
==================================
Max supported route count   : 131072
Active route count: 8

-------------------------------------------
 VNID     Total    MACONLY  MACIPv4  MACIPv6
-------------------------------------------
 10       8        4        2        2

Total number of entries are 1
```

show evpn srv6 static host state

Use this command to display the state of the host which is configured statically.

Command Syntax

```
show evpn srv6 static host state
```

Parameters

None

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates to display the show output of `evpn srv6 static host status`

```
PPE1#show evpn srv6 static host status
SRv6 Static Host Information
============================
   Codes: NW - Network Port
          AC - Access Port
         (u) - Untagged

VNID    Ifname      Outer-Vlan Inner-vlan Ip-Addr
Mac-Addr       Status
_____________________________________________________________________________
_____________________________
10      xe29.100    ---        ---        0.0.0.0
0020.9400.0003 Active
10      xe29.100    ---        ---        2001::9
0030.9400.0003 Inactive

Total number of entries are 2
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Ethernet VPN (EVPN)	A solution that provides Ethernet multipoint services over MPLS networks, enabling control-plane-based MAC learning in the core.

Virtual Private LAN Service (VPLS)	An early MPLS VPN technology that provides multipoint-to-multipoint wide-area Ethernet services for enterprise users.
MP-BGP Protocol	Multi-Protocol Border Gateway Protocol, used for control-plane MAC learning in EVPN instances.
Control Plane	The part of a network responsible for routing protocols, forwarding tables, and other control functions.
Data Plane	The part of a network responsible for forwarding user data based on the information in the control plane.
Route Reflector (RR)	A device in a network that helps to reduce the number of IBGP connections required in a full-mesh topology by reflecting routes from one IBGP speaker to another.
Media Access Control (MAC) Address	A unique identifier assigned to network interfaces for communication at the data link layer of a network segment.
BGP Extensions	Additional functionality added to the Border Gateway Protocol (BGP) to support specific requirements or features.
IMET Route	A route type in EVPN used for Broadcast, Unknown Unicast, and Multicast (BUM) traffic delivery across EVPN networks.
Ethernet Segment Route	A route type in EVPN used in multi-homing scenarios and for Designated Forwarder Election.
Ingress Replication (IR)	A technique used in multicast routing to replicate multicast traffic at the ingress router and forward it to multiple egress routers.
Designated Forwarder (DF)	In EVPN, the PE responsible for sending broadcast, unknown multicast, and multicast (BUM) traffic to the CE on a particular Ethernet Segment.

CHAPTER 5 BGP ORF Prefix-List VPNV4 Address

Overview

The Border Gateway Protocol (BGP) Outbound Route Filtering feature operates as a Prefix-Based filtering system within the BGP. Its primary purpose is to reduce the volume of BGP updates exchanged among peer routers. By selectively screening out unnecessary routing updates at the source, this feature effectively lessens the strain on resources needed for generating and handling routing updates. Its objective is to streamline router processing, especially for routers not set up to accept full BGP route updates from a service provider network.

Feature Characteristics

This feature provides customers with various routing options, such as access to routing information like a full table view, solely a default route, or a tailored subset such as a default route combined with locally originated prefixes from the service provider. Typically, BGP service providers do not impose complex outbound filtering policies on their customers.

Benefits

The advantages of Prefix-Based Outbound Route Filtering:

- Minimize unnecessary routing updates
- Reduces resources required for routing update generation and processing
- Reduces required to receive and discard routes.

Configuration

The BGP Prefix-Based Outbound Route Filtering feature offers support for prefix length matching, wildcard-based prefix matching, and exact address prefix matching across various address families. It allows configuration on a router to enable Outbound Route Filtering (ORF) capabilities for sending or receiving, using the "send" or "receive" keywords. Moreover, it permits configuration to enable both sending and receiving ORF capabilities using the "both" keyword.

Topology

In this topology, the PE1, P1, RR1, and PE2 interface is established. It allows configuration on a router to enable Outbound Route Filtering (ORF) capabilities.

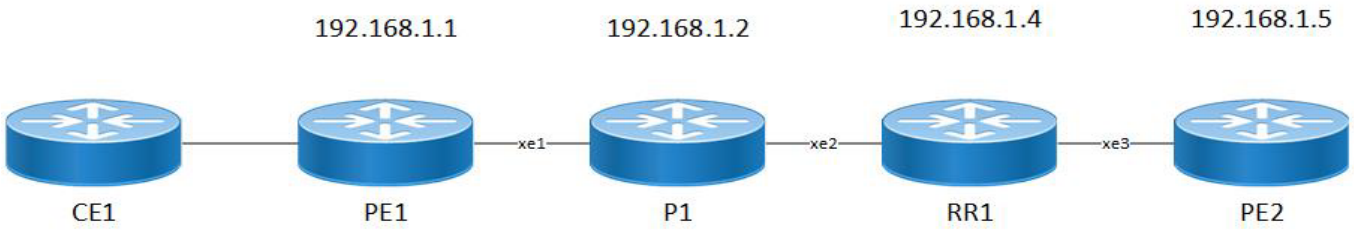

ORF-Prefix VPNv4 Address Topology

PE1

The Provider Edge (PE) 1 is a device at the edge of a service provider's network. For the PE1 configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.
```
PE1(config)#interface lo
PE1(config-if)#ip address 192.168.1.1/32 secondary
PE1(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.
```
PE1(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for PE-1. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.
```
PE1(config)#router ldp
PE1(config-router)#router-id 192.168.1.1
PE1(config-router)#targeted-peer ipv4 192.168.1.5
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#transport-address ipv4 192.168.1.1
```

4. To exit from the router mode for LDP, execute the following command.
```
PE1(config)#exit
```

5. In interface mode, assign an IP address to the interface. Then, activate label switching capability on the interface and enable LDP on it.
```
PE1(config)#interface xe1
PE1(config-if)#ip address 12.1.1.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)# ip ospf cost 10
```

6. To exit from the interface configuration on network side, execute the following command.
```
PE1(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.
```
PE1(config)# router ospf 100
```

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).
```
PE1(config-router)# ospf router-id 192.168.1.1
PE1(config-router)# bfd all-interfaces
PE1(config-router)# timers spf exp 50 50
PE1(config-router)# timers throttle lsa all 0 1 1
PE1(config-router)# network 12.1.1.0/24 area 0.0.0.0
PE1(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.
```
PE1(config)#exit
```

10. Switch to BGP router mode. Establish PE1 as an iBGP peer. Specify the loopback as the source for iBGP peering with the remote PE1 router. Activate PE1 in the VPNv4 unicast address family.
```
PE1(config)#router bgp 100
PE1(config-router)# bgp router-id 192.168.1.1
PE1(config-router)# neighbor 192.168.1.4 remote-as 100
PE1(config-router)# neighbor 192.168.1.4 update-source lo
PE1(config-router)# neighbor 192.168.1.4 advertisement-interval 0
PE1(config-router)# address-family vpnv4 unicast
PE1(config-router-af)# neighbor 192.168.1.4 activate
PE1(config-router-af)# neighbor 192.168.1.4 capability orf prefix-list receive
PE1(config-router-af)# exit-address-family
```

11. To exit from the BGP, execute the following command.
```
PE1(config)#exit
```

P

The Provider (P) is a device at the edge of a service provider's network. For the P configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.
```
P(config)#interface lo
P(config-if)#ip address 192.168.1.1/32 secondary
P(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.
```
P(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for P. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.
```
P(config)#router ldp
P(config-router)#router-id 192.168.1.2
P(config-router)#targeted-peer ipv4 192.168.1.2
```

4. To exit from the router mode for LDP, execute the following command.
```
P(config)#exit
```

5. In interface mode, assign an IP address to the interface. Then, activate label switching capability on the interface and enable LDP on it.
```
P(config)#interface xe1
P(config-if)#ip address 12.1.1.1/24
P(config-if)#label-switching
P(config)#interface xe1
P(config-if)#ip address 13.1.1.2/24
P(config-if)#label-switching
```

6. To exit from the interface configuration on network side, execute the following command.
```
P(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.
```
P(config)# router ospf 100
```

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).
```
P(config-router)# ospf router-id 192.168.1.2
P(config-router)# bfd all-interfaces
P(config-router)# timers spf exp 50 50
P(config-router)# timers throttle lsa all 0 1 1
P(config-router)# network 12.1.1.0/24 area 0.0.0.0
P(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.
```
P(config)#exit
```

RR

Route Reflector (RR) is a designated router that will reflect routes learned from other iBGP peers. All routers form a peering relationship only with the Route Reflector. For the RR configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.
```
RR(config)#interface lo
RR(config-if)#ip address 192.168.1.1/32 secondary
RR(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.
```
RR(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for P. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.
```
RR(config)#router ldp
RR(config-router)#router-id 192.168.1.4/32
RR(config-router)#transport-address ipv4 192.168.1.4
```

4. To exit from the router mode for LDP, execute the following command.
```
RR(config)#exit
```

5. In interface mode, set up the IP address for the interface and activate label switching capability on it.
```
RR(config)#interface xe1
RR(config-if)#ip address 12.1.1.1/24
RR(config-if)#label-switching
RR(config-if)#enable-ldp ipv4
RR(config)#interface xe3
RR(config-if)#ip address 14.1.1.1/24
RR(config-if)#label-switching
RR(config-if)#enable-ldp ipv4
```

6. To exit from the interface configuration on network side, execute the following command.
```
RR(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.
```
RR(config)# router ospf 100
```

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).
```
RR(config-router)# ospf router-id 192.168.1.4
RR(config-router)# bfd all-interfaces
RR(config-router)# timers spf exp 50 50
RR(config-router)# timers throttle lsa all 0 1 1
RR(config-router)# network 12.1.1.0/24 area 0.0.0.0
RR(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.
```
RR(config)#exit
```

10. Switch to BGP router mode. Establish RR as an iBGP peer. Specify the loopback as the source for iBGP peering with the remote RR router. Activate RR in the VPNv4 unicast address family.
```
RR(config)#router bgp 100
RR(config-router)# bgp router-id 192.168.1.1
RR(config-router)# neighbor 192.168.1.4 remote-as 100
RR(config-router)# neighbor 192.168.1.4 update-source lo
RR(config-router)# neighbor 192.168.1.4 advertisement-interval 0
RR(config-router)# address-family vpnv4 unicast
RR(config-router-af)# neighbor 192.168.1.4 active
RR(config-router-af)# neighbor 192.168.1.4 capability orf prefix-list s
RR(config-router-af)# neighbor 192.168.1.1 prefix-list
RR(config-router-af)# exit-address-family
```

11. To exit from the BGP, execute the following command.
```
PE1(config)#exit
```

12. To configure the global prefix, execute the following command in the global mode.
```
RR(config)# ip prefix-list ORF1
RR(config-ip-prefix-list)# seq 1 permit 45.1.1.0/24
```

13. To exit from the BGP, execute the following command.
```
RR(config)#exit
```

PE2

The Provider Edge (PE) 2 is a device at the edge of a service provider's network. For the PE2 configuration, follow these steps.

1. Enable activate the loopback interface, enter the following command while in configuration mode. Then, proceed to set up the IP address for the loopback interface.
```
PE2(config)#interface lo
PE2(config-if)#ip address 192.168.1.1/32 secondary
PE2(config-if)#enable-ldp ipv4
```

2. To exit from the loopback interface, execute the following command.
```
PE2(config)#exit
```

3. Set up the Router ID. Next, configure targeted LDP sessions for P. Once done, exit from targeted-peer mode. Then, configure the transport address for LDP to run on an IPv4 interface for TCP sessions.
```
PE2(config)#router ldp
PE2(config-router)#router-id 192.168.1.1
PE2(config-router)#targeted-peer ipv4 192.168.1.5
PE2(config-router-targeted-peer)#exit-targeted-peer-mode
PE2(config-router)#transport-address ipv4 192.168.1.1
PE2(config-router)#transport-address ipv4 192.168.1.5
```

4. To exit from the router mode for LDP, execute the following command.
```
PE2(config)#exit
```

5. In interface mode, assign an IP address to the interface. Then, activate label switching capability on the interface and enable LDP on it.
```
PE2(config)#interface xe1
PE2(config-if)#ip address 12.1.1.1/24
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)# ip ospf cost 10
```

6. To exit from the interface configuration on network side, execute the following command.
```
PE2(config)#exit
```

7. Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process.
```
PE2(config)# router ospf 100
```

8. Configure the OSPF Router ID and define the interface for OSPF operation and link it with the area ID (0).
```
PE2(config-router)# ospf router-id 192.168.1.1
PE2(config-router)# bfd all-interfaces
PE2(config-router)# timers spf exp 50 50
PE2(config-router)# timers throttle lsa all 0 1 1
PE2(config-router)# network 12.1.1.0/24 area 0.0.0.0
PE2(config-router)# network 192.168.1.1/32 area 0.0.0.0
```

9. To exit from the OSPF, execute the following command.
```
PE2(config)#exit
```

10. Switch to BGP router mode. Establish PE2 as an iBGP peer. Specify the loopback as the source for iBGP peering with the remote PE2 router. Activate PE2 in the VPNv4 unicast address family.
```
PE2(config)#router bgp 100
PE2(config-router)# bgp router-id 192.168.1.1
PE2(config-router)# neighbor 192.168.1.4 remote-as 100
PE2(config-router)# neighbor 192.168.1.4 update-source lo
PE2(config-router)# neighbor 192.168.1.4 advertisement-interval 0
PE2(config-router)# address-family vpnv4 unicast
PE2(config-router-af)# neighbor 192.168.1.4 activate
PE2(config-router-af)# exit-address-family
```

11. To exit from the BGP, execute the following command.
```
PE2(config)#exit
```

Validation

```
PE1-7017#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid,
 > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop            Metric    LocPrf     Weight Path
Route Distinguisher: 192.168.1.1:100 (Default for VRF vrf100)
*> l 45.1.1.0/24      100.1.1.2             0        100        0   400 i
*> l 45.1.2.0/24      100.1.1.2             0        100        0   400 i
```

```
*> l 45.1.3.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.4.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.5.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.6.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.7.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.8.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.9.0/24      100.1.1.2            0        100       0   400 i
*> l 45.1.10.0/24     100.1.1.2            0        100       0   400 i
*> l 100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:101 (Default for VRF vrf101)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:102 (Default for VRF vrf102)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
```

```
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.3.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:103 (Default for VRF vrf103)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
*>   100.1.1.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0              0        100       32768  ?
*> l 100.1.4.0/24     0.0.0.0              0        100       32768  ?
*>   100.1.5.0/24     0.0.0.0              0        100       32768  ?
*>il 200.1.1.0        192.168.1.4          0        100       0    ?
*>il 200.1.2.0        192.168.1.4          0        100       0    ?
*>il 200.1.3.0        192.168.1.4          0        100       0    ?
*>il 200.1.4.0        192.168.1.4          0        100       0    ?
*>il 200.1.5.0        192.168.1.4          0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.1:104 (Default for VRF vrf104)
*>   45.1.1.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.2.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.3.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.4.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.5.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.6.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.7.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.8.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.9.0/24      100.1.1.2            0        100       0   400 i
*>   45.1.10.0/24     100.1.1.2            0        100       0   400 i
```

```
*>   100.1.1.0/24     0.0.0.0             0        100       32768  ?
*>   100.1.2.0/24     0.0.0.0             0        100       32768  ?
*>   100.1.3.0/24     0.0.0.0             0        100       32768  ?
*>   100.1.4.0/24     0.0.0.0             0        100       32768  ?
*> l 100.1.5.0/24     0.0.0.0             0        100       32768  ?
*>il 200.1.1.0        192.168.1.4         0        100       0    ?
*>il 200.1.2.0        192.168.1.4         0        100       0    ?
*>il 200.1.3.0        192.168.1.4         0        100       0    ?
*>il 200.1.4.0        192.168.1.4         0        100       0    ?
*>il 200.1.5.0        192.168.1.4         0        100       0    ?
 Announced routes count = 15
 Accepted routes count = 5
Route Distinguisher: 192.168.1.6:100
*>il 200.1.1.0        192.168.1.4         0        100       0    ?
* il                  192.168.1.5         0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:101
*>il 200.1.2.0        192.168.1.4         0        100       0    ?
* il                  192.168.1.5         0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:102
*>il 200.1.3.0        192.168.1.4         0        100       0    ?
* il                  192.168.1.5         0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:103
*>il 200.1.4.0        192.168.1.4         0        100       0    ?
* il                  192.168.1.5         0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2
Route Distinguisher: 192.168.1.6:104
*>il 200.1.5.0        192.168.1.4         0        100       0    ?
* il                  192.168.1.5         0        100       0    ?
 Announced routes count = 0
 Accepted routes count = 2x
```

Use these commands to validate the BGP Neighbor Table.

```
PE1#show ip bgp neighbors 192.168.1.4
BGP neighbor is 192.168.1.4, remote AS 100, local AS 100, internal link, peer in
dex: 32
  BGP version 4, local router ID 192.168.1.1, remote router ID 192.168.1.4
  BGP state = Established, up for 00:05:03
  Last read 00:00:18, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Labeled-Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
```

```
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 3229 messages, 1 notifications, 0 in queue
  Sent 3252 messages, 2 notifications, 0 in queue
  Route refresh request: received 2, sent 0
  Minimum time between advertisement runs is 0 seconds
  Update source is lo

For address family: IPv4 Unicast  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  1 announced prefixes


For address family: VPNv4 Unicast  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (64) Prefix-list:
      Send-mode: received
      Receive-mode: advertised
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode: received
      Receive-mode: advertised
  Outbound Route Filter (ORF): received (1 entries)
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  1 announced prefixes

For address family: IPv4 Labeled-Unicast  BGP table version 6, neighbor version
5
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  1 announced prefixes

For address family: L2VPN VPLS  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
```

```
  9 announced prefixes

For address family: L2VPN EVPN  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  20 accepted prefixes
  Accepted AD:10 MACIP:0 MCAST:10 ESI:0 PREFIX:0
  20 announced prefixes

For address family: IPv6 Unicast  BGP table version 2, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

For address family: VPNv6 Unicast  BGP table version 2, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  4 accepted prefixes
  5 announced prefixes

For address family: 6PE Labeled Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  10 accepted prefixes
  10 announced prefixes

 Connections established 4; dropped 3
Local host: 192.168.1.1, Local port: 179
Foreign host: 192.168.1.4, Foreign port: 44897
TCP MSS: (0), Advertise TCP MSS: (9176), Send TCP MSS: (1460),  Receive TCP MSS:
 (1460)
Sock FD : (43)
Nexthop: 192.168.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:03, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)
```

Use these commands to validate the RR configuration.

```
# RR1-7038#SH IP BGP VPNV4 ALL
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid,
 > best, i - internal, l - labeled
              S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete

     Network          Next Hop            Metric    LocPrf     Weight Path
Route Distinguisher: 192.168.1.1:100
*>il 45.1.1.0/24      192.168.1.1           0         100        0   400 i
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:100
*>il 200.1.1.0        192.168.1.6           0         100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:101
*>il 200.1.2.0        192.168.1.6           0         100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:102
*>il 200.1.3.0        192.168.1.6           0         100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:103
*>il 200.1.4.0        192.168.1.6           0         100        0    ?
 Announced routes count = 0
 Accepted routes count = 1
Route Distinguisher: 192.168.1.6:104
*>il 200.1.5.0        192.168.1.6           0         100        0    ?
 Announced routes count = 0
 Accepted routes count = 1

RR1-7038#show ip bgp neighbors 192.168.1.1
BGP neighbor is 192.168.1.1, remote AS 100, local AS 100, internal link, peer in
dex: 2
  BGP version 4, local router ID 192.168.1.4, remote router ID 192.168.1.1
  BGP state = Established, up for 00:06:19
  Last read 00:00:09, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv4 Labeled-Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 3244 messages, 2 notifications, 0 in queue
  Sent 3242 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 2
  Minimum time between advertisement runs is 0 seconds
  Update source is lo
```

```
For address family: IPv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  4 announced prefixes

For address family: VPNv4 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  AF-dependant capabilities:
    Outbound Route Filter (ORF) type (64) Prefix-list:
      Send-mode: advertised
      Receive-mode: received
    Outbound Route Filter (ORF) type (128) Prefix-list:
      Send-mode: advertised
      Receive-mode: received
  Outbound Route Filter (ORF): sent;
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  Inbound path policy configured
  Incoming update prefix filter list is *ORF1
  1 accepted prefixes
  5 announced prefixes

For address family: IPv4 Labeled-Unicast  BGP table version 6, neighbor version
6
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  NEXT_HOP is always this router
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  4 announced prefixes

For address family: L2VPN VPLS  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  9 announced prefixes

For address family: L2VPN EVPN  BGP table version 5, neighbor version 5
  Index 1, Offset 0, Mask 0x2
```

```
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  20 accepted prefixes
  Accepted AD:10 MACIP:0 MCAST:10 ESI:0 PREFIX:0
  20 announced prefixes

For address family: IPv6 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

For address family: VPNv6 Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  5 accepted prefixes
  5 announced prefixes

For address family: 6PE Labeled Unicast  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Route-Reflector Client
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  10 accepted prefixes
  10 announced prefixes

Connections established 4; dropped 3
Local host: 192.168.1.4, Local port: 44897
Foreign host: 192.168.1.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS:
 (1460)
Sock FD : (26)
Nexthop: 192.168.1.4
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:06:19, due to BGP Notification received
Notification Error Message: (Cease/Administratively Reset.)

RR1-7038#
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
ORF	ORF stands for Outbound Route Filtering. It is a feature in routing protocols, particularly in BGP (Border Gateway Protocol), that allows a router to advertise to its neighbor routers the set of routes it can accept or reject.
LDP	LDP stands for Label Distribution Protocol. It is a signaling protocol used in MPLS (Multiprotocol Label Switching) networks to distribute and exchange labels between MPLS-enabled routers.
OSPF	OSPF stands for Open Shortest Path First. It is a routing protocol used in computer networks, particularly in large enterprise and service provider networks.

Improved Routing

This section describes the new features for improved network routing introduced in the Release 6.5.3.

Release 6.5.3

- Segment Routing ECMP for ISIS or OSPF

Release 6.5.2

- ISIS Multi Topology

CHAPTER 1 Segment Routing ECMP for ISIS or OSPF

Overview

Segment Routing (SR) is a source-based routing technique where you can specify a route in a network through which a packet is sent. The path that a particular packet needs to traverse is represented by one or more segments (nodes and links).

Equal Cost Multipath (ECMP) refers to single-hop, equal-cost links between adjacent nodes with a forwarding mechanism for routing traffic along multiple paths of equal cost. For ECMP enabled devices, OcNOS uses Forwarding Plane Load Balancing and installs the maximum number of ECMP routes supported by the kernel. This allows for load balancing to be performed with more than one next-hop to reach a destination.

SR with ECMP support for Intermediate System to Intermediate System (ISIS) and Open Shortest Path First (OSPF), selects all the valid equal-cost next-hop peers of an IP prefix and creates ECMP Incoming Label Maps (ILM) and FEC-to-NHLFE (FTN) entries for that prefix with all the IS-IS/OSPF SR next-hops in the forwarding plane.

Feature Characteristics

The main characteristics of SR ECMP are as follows:

- Distributes packets across multiple logical paths (LSP) carrying qualified traffic over MPLS underlay using SR as a transport. The traffic is distributed based on a collection of such LSPs, known as an ECMP set.
- Uses an internal hashing algorithm by the forwarding plane to distribute traffic among multiple next-hops, assigning the traffic flow to a particular next-hop.
- When TI-LFA is enabled, IGP adds the ECMP next-hops as primary and computes and adds a backup for each of the ECMP next-hops in the FTN and ILM entry of the prefix.
- When TI-LFA is disabled, IGP computes and adds all the ECMP next-hops in the FTN and ILM entry of the prefix.

Note: Load balancing on ECMP next-hops does not guarantee equal distribution of traffic across the ECMP paths. Load balancing in the hardware is done using hashing of a combination of headers in the traffic streams, such as src, dst mac, ip pair and so on. The unique combination of such headers may result in the same hash which in turn results in the same ECMP next-hop. This causes unequal distribution of traffic within the ECMP next-hop interfaces.

Benefits

The key benefits of SR ECMP are as follows:

- Distributes traffic across multiple equal-cost paths, effectively balancing the load, optimizing resource utilization throughout the network, and preventing congestion.
- Reroutes traffic to alternative equal-cost paths in case of a link or node failure, thus reducing downtime and maintaining continuous service.
- Offers redundancy by utilizing multiple paths such that if one path becomes unavailable, traffic is redirected to other paths seamlessly, bolstering network resilience and reliability, when TI-LFA is enabled.

Prerequisites

The SR ECMP feature can be enabled on the following devices:

- OcNOS devices that support ISIS/OSPF Segment Routing.
- OcNOS devices that support MPLS services such as VPLS,VPWS,L3VPN,6PE,6VPE and EVPN (ELINE,ELAN,ETREE).

Configuration

The following configuration enables ECMP with ISIS-SR for L3VPN and EVPN ELINE services.

Topology

This topology includes Edge nodes - DUT1 and DUT2, Intermediate nodes - DUT2, DUT3, DUT4, and DUT5 and Target Groups 1 and 2.

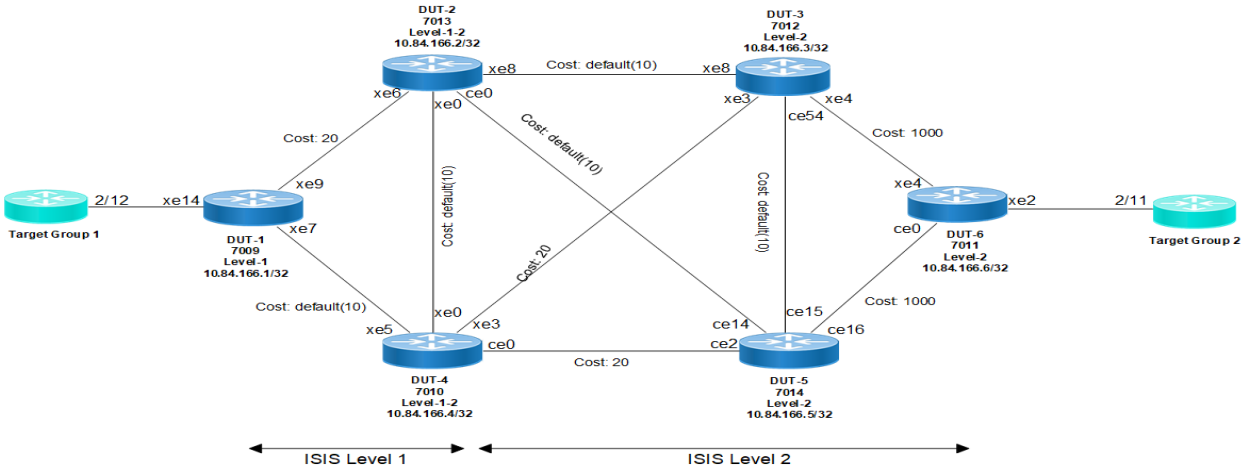

The ECMP Label Switched Path (LSP) derived from the above topology is as follows:

Source	Destination	ECMP	Path	Cost
DUT1	DUT2	YES	DUT1-DUT2	30
			DUT1-DUT4-DUT2	30
DUT1	DUT6	YES	DUT1-DUT2-DUT3-DUT6	1040
			DUT1-DUT4-DUT5-DUT6	1040

To configure SR ECMP functionality on PE nodes with ISIS as IGP, follow the steps mentioned below:

1. Configure loop-back interface.
    1. Access interface configuration mode for the loopback interface (`interface lo`).
    2. Assign an IPv4 address to the loopback interface using the IPv4 address command followed by the desired IPv4 address and subnet mask (`ipv6 address 10.84.166.1/32`).

3. Assign appropriate prefix-sid index for the loopback interface (`prefix-sid index 100 no-php`).
4. Configure IS-IS for IPv4 on the loopback interface using the ip router isis command, specifying the IS-IS process ID (`ip router isis 1`).

```
DUT1(config)#interface lo
DUT1(config-if)# ip address 127.0.0.1/8
DUT1(config-if)# ip address 10.84.166.1/32 secondary
DUT1(config-if)# ipv6 address ::1/128
DUT1(config-if)# prefix-sid index 100 no-php
DUT1(config-if)# ip router isis 1
DUT1(config-if)# exit
```

2. Configure network interface.
   1. Access interface configuration mode for the desired network interface (`interface xe9 and xe7`).
   2. Assign an IPv4 address to the loopback interface using the ipv4 address command followed by the desired IPv4 address and subnet mask (`ip address 10.11.22.1/30`).
   3. Configure the MTU for the interface (`mtu 9216`).
   4. Configure IS-IS for IPv4 on the interface using the IP router ISIS command, specifying the IS-IS process ID (`ip router isis 1`).

```
DUT1(config)#interface xe9
DUT1(config-if)# load-interval 30
DUT1(config-if)# ip address 10.11.22.1/30
DUT1(config-if)# mtu 9216
DUT1(config-if)# label-switching
DUT1(config-if)# ip router isis 1
DUT1(config-if)#
DUT1(config-if)#exit
```

3. In global configuration mode, perform the following as shown in the configuration snapshots below:
   1. Configure ISIS Settings
   2. Perform the BGP Configuration
   3. Create IP VRF:
   4. Define the L3VPN access intf configuration and IP VRF mapping.
   5. Create MAC VRF.
   6. Define the ELINE instance and with the MAC VRF Mapping and access intf configuration
   7. Enable ECMP for SR entities for FTN as its PE Edge node using the command `mpls ftn-ecmp sr`

      Note: Use ECMP SR for ILM in case of P transit nodes.

Configuration Snapshot

Edge Nodes (DUT1 and DUT6)

```
DUT1#sh run
!
!Software version: UFI_S9510-30XC-OcNOS-SP-PLUS-6.6.0.99-Alpha 10/
  07/2024 21:37:20
```

```
!
! Last configuration change at 00:10:29 UTC Thu Nov 16 2023 by root
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
logging console 5
logging level all 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
mpls ilm-ecmp sr
mpls ftn-ecmp sr
!
hostname DUT1
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
!
evpn mpls enable
!
```

```
evpn mpls irb
!
ip vrf management
!
ip vrf vrf701
 rd 10:701
 route-target both 10:701
!
mac vrf ELINE_DUT1_DUT6_501
 rd 10.84.166.1:501
 route-target both 501:501
!
evpn mpls vtep-ip-global 10.84.166.1
!
evpn mpls id 501 xconnect target-mpls-id 1501
 host-reachability-protocol evpn-bgp ELINE_DUT1_DUT6_501
!
router ldp
 targeted-peer ipv4 10.84.166.6
  exit-targeted-peer-mode
!
interface ce0
!
interface ce1
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.84.166.1/32 secondary
 ipv6 address ::1/128
 prefix-sid index 100 no-php
 ip router isis 1
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
```

```
 ipv6 address ::1/128
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
 speed 10g
 load-interval 30
 ip address 10.11.44.1/30
 mtu 9216
 label-switching
 ip ospf network point-to-point
 ip router isis 1
!
interface xe8
!
interface xe9
 load-interval 30
 ip address 10.11.22.1/30
 mtu 9216
 label-switching
 ip router isis 1
 isis wide-metric 20
!
interface xe10
!
interface xe11
!
interface xe12
!
interface xe13
!
```

```
interface xe14
 mtu 9216
!
interface xe14.501 switchport
 description ELINE_DUT1_DUT6_501
 encapsulation dot1q 501
 load-interval 30
 mtu 9216
 access-if-evpn
  map vpn-id 501
!
interface xe14.701
 encapsulation dot1q 701
 load-interval 30
 ip vrf forwarding vrf701
 ip address 100.7.1.1/24
 mtu 9216
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
interface xe21
!
interface xe22
!
interface xe23
!
interface xe24
!
interface xe25
```

```
!
interface xe26
!
interface xe27
!
interface xe28
!
interface xe29
!
 exit
!
router isis 1
 is-type level-1
 metric-style wide
 mpls traffic-eng router-id 10.84.166.1
 mpls traffic-eng level-1
 capability cspf
 bfd all-interfaces
 net 49.0001.0000.0001.0011.00
 isis segment-routing global block 20000 23000
 segment-routing mpls
!
router bgp 65010
 bgp router-id 10.84.166.1
 neighbor 10.84.166.6 remote-as 65010
 neighbor 10.84.166.6 update-source lo
 !
 address-family vpnv4 unicast
 neighbor 10.84.166.6 activate
 exit-address-family
 !
 address-family l2vpn evpn
 neighbor 10.84.166.6 activate
 exit-address-family
 !
 address-family ipv4 vrf vrf701
 redistribute connected
 neighbor 100.7.1.2 remote-as 101
 neighbor 100.7.1.2 activate
```

```
 exit-address-family
 !
 exit
!
!
end


DUT1#
DUT1#
```

Transit Nodes (DUT2, DUT3, DUT4, and DUT5)

```
DUT2#
DUT2#sh run
!
!Software version: EC_AS5916-54X-OcNOS-SP-MPLS-6.5.3.86-Alpha 10/1
3/2024 14:39:27
!
!Last configuration change at 13:12:58 UTC Mon Oct 14 2024 by ocno
s
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
logging console 5
logging level all 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
!
mpls ilm-ecmp sr
!
hostname DUT2
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
!
ip vrf management
```

```
!
interface ce0
 load-interval 30
 ip address 10.22.55.1/30
 mtu 9216
 label-switching
 ip ospf network point-to-point
 ip router isis 1
!
interface ce1
!
interface ce2
!
interface ce3
!
interface ce4
!
interface ce5
!
interface eth0
 ip vrf forwarding management
 ip address 192.168.3.10/24
!
interface lo
 ip address 127.0.0.1/8
 ip address 10.84.166.2/32 secondary
 ipv6 address ::1/128
 prefix-sid index 200 no-php
 ip router isis 1
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe0
 load-interval 30
 ip address 10.22.44.1/30
 mtu 9216
 label-switching
 ip ospf network point-to-point
 ip router isis 1
!
interface xe1
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
 load-interval 30
 ip address 10.11.22.2/30
 mtu 9216
```

```
 label-switching
 ip ospf network point-to-point
 ip router isis 1
 isis wide-metric 20
!
interface xe7
!
interface xe8
 load-interval 30
 ip address 10.22.33.1/30
 mtu 9216
 label-switching
 ip ospf network point-to-point
 ip router isis 1
!
interface xe9
!
interface xe10
!
interface xe11
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
interface xe21
!
interface xe22
!
interface xe23
!
interface xe24
!
interface xe25
!
interface xe26
!
interface xe27
!
interface xe28
!
interface xe29
!
```

```
interface xe30
!
interface xe31
!
interface xe32
!
interface xe33
!
interface xe34
!
interface xe35
!
interface xe36
!
interface xe37
!
interface xe38
!
interface xe39
!
interface xe40
!
interface xe41
!
interface xe42
!
interface xe43
!
interface xe44
!
interface xe45
!
interface xe46
!
interface xe47
!
 exit
!
router isis 1
 is-type level-1-2
 metric-style wide
 mpls traffic-eng router-id 10.84.166.2
 mpls traffic-eng level-1
 mpls traffic-eng level-2
 capability cspf
 bfd all-interfaces
 net 49.0001.0000.0001.0022.00
 redistribute isis level-2 into level-1
 isis segment-routing global block 20000 23000
 segment-routing mpls
!
ip route vrf management 0.0.0.0/0 192.168.3.1 eth0
!
!
end

DUT2#
```

Validation

Validation of SR-ECMP on DUT1[Edge Router]

Here are the show outputs that display the ISISv4 neighbour and routing information with ECMP for DUT1.

```
DUT1#sh clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0001.0044 xe7         b86a.97c8.3dcb      Up     19        L1   IS-IS
0000.0001.0022 xe9         80a2.352b.7008      Up     19        L1   IS-IS
DUT1#
DUT1#sh clns neighbors detail

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
0000.0001.0044 xe7         b86a.97c8.3dcb      Up     21        L1   IS-IS
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 00:04:27
  Area Address(es): 49.0001
  IP Address(es): 10.11.44.2
  Level-1 Protocols Supported: IPv4
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
  Adjacency SID: 26880, ILM ID: 3


0000.0001.0022 xe9         80a2.352b.7008      Up     21        L1   IS-IS
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 00:04:27
  Area Address(es): 49.0001
  IP Address(es): 10.11.22.2
  Level-1 Protocols Supported: IPv4
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
  Adjacency SID: 26881, ILM ID: 4


DUT1#
```

```
DUT1#sh ip route 10.84.166.6/32
VRF: Default, Routing entry for 10.84.166.6/32
  Known via "isis", distance 115, metric 1040,  External Route Tag: 0, installed
00:22:23, best
  Last update 00:22:23 ago
  * 10.11.22.2, via xe9
  * 10.11.44.2, via xe7

DUT1#
```

The following show outputs displays the validation for L3VPN.

```
DUT1#sh ip bgp vpnv4 all neighbors
BGP neighbor is 10.84.166.6, remote AS 65010, local AS 65010, internal link, peer index:
4
  BGP version 4, local router ID 10.84.166.1, remote router ID 10.84.166.6
  BGP state = Established, up for 01:18:17
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN EVPN: advertised and received
  Received 351 messages, 0 notifications, 0 in queue
  Sent 333 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

For address family: VPNv4 Unicast  BGP table version 9, neighbor version 9
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes


For address family: L2VPN EVPN  BGP table version 6, neighbor version 6
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  Accepted AD:1 MACIP:0 MCAST:0 ESI:0 PREFIX:0
  1 announced prefixes

 Connections established 3; dropped 2
Local host: 10.84.166.1, Local port: 179
Foreign host: 10.84.166.6, Foreign port: 40371
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (536)
Sock FD : (29)
```

```
Nexthop: 10.84.166.1 lo
Nexthop global: :: lo
Nexthop local: :: lo
BGP connection: non shared networkLast Reset: 01:19:41, due to Hold Timer Expired
(Notification sent)
Notification Error Message: (Hold Timer Expired/No sub-error code)

DUT1#sh mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP or SR ECMP

Code    FEC               FTN-ID VRF-ID    Nhlfe-ID    Pri   Out-Label    Out-Intf
Nexthop         UpTime
   B>   200.7.1.0/24       1        2      51          Yes   25600        -
10.84.166.6     00:05:18
DUT1#
```

The following show output displays the validation for EVPN ELINE.

```
DUT1#show bgp l2vpn evpn summary
BGP router identifier 10.84.166.1, local AS number 65010
BGP table version is 6
1 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/
PfxRcd     AD  MACIP  MCAST    ESI  PREFIX-ROUTE
10.84.166.6       4 65010       353        336      6     0      0  01:19:10
1      1      0      0      0         0

Total number of neighbors 1

Total number of Established sessions 1
DUT1#sh evpn mpls xconnect
EVPN Xconnect Info
========================
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                              Remote       Connection-Details
================================== ============
================================================================================
VPN-ID       EVI-Name      MTU   VPN-ID       Source       Destination
PE-IP           MTU   Type   NW-Status
================================== ============
================================================================================
```

```
501          ----          9216  1501         xe14.501     --- Single Homed Port ---
10.84.166.6     9216  AC-NW  NW-SET

Total number of entries are 1

DUT1#sh evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels
 (*) in Policy - tunnel-policy inherited from mac-vrf
 (e) - Service mapped over MPLS Multipath/ECMP
================+===========+=========+=========+===============+=========+=========+==
==========+=========+============+=========
                            Local     Remote                      Local     Remote
MPLS-Multipath          Underlay
 Destination     Status      VPWS-ID   VPWS-ID   Policy          UC-Label  UC-Label  Grp-
Name     NHLFE-ix  NW-Intf      NW-Label
================+===========+=========+=========+===============+=========+=========+==
==========+=========+============+=========
 10.84.166.6    Installed   501      1501     --            26240    26240    --
40        NA(e)         NA(e)

Total number of entries are 1
DUT1#
DUT1#sh evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status        Up/Down       Update        local-evpn-
id remote-evpn-id
================================================================================
=================
10.84.166.1      10.84.166.6      Installed     00:08:06      00:08:06      501
1501

Total number of entries are 1
DUT1#
```

The following show output displays the ECMP validation for ISIS-SR.

```
This command displays the ILM-ID, FTN-ID, In-Label & Out-Label for all the IS-IS routes
which have ILM/FTN entry installed. For the ECMP prefix 10.84.166.6//32, only one ILM/
FTN entry will be installed, but Out-Label will be separate for each nexthop.2

DUT1#sh ip isis route prefix 10.84.166.6/32 detail

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      ILM-ID     FTN-ID     In-Label   Next-Hop
Interface      Tag       Out-Label
ia   10.84.166.6/32      1040        9          5          20600      10.11.44.2
xe7            0         20600
                                                          10.11.22.2        xe9
0         20600
```

```
     Src: 0000.0001.0022 Ifindex 10016
     Src: 0000.0001.0044 Ifindex 10014
DUT1#
```

This command displays In-Label and Out-Label of all next-hops of the FEC.

```
DUT1#sh isis segment-routing label detail
Tag 1  Segment-Routing: Label Table
FEC                In-Label       Out-Label      Out-Intf    Nexthop           Dependent
Tunnels
10.84.166.4/32      20400          20400          xe7         10.11.44.2
10.84.166.2/32      20200          20200          xe7         10.11.44.2
                                   20200          xe9         10.11.22.2
10.84.166.1/32      20100          N/A            lo          127.0.0.1
10.84.166.3/32      20300          20300          xe7         10.11.44.2
                                   20300          xe9         10.11.22.2
10.11.44.2/32       26880          3              xe7         10.11.44.2
10.84.166.5/32      20500          20500          xe7         10.11.44.2
                                   20500          xe9         10.11.22.2
10.84.166.6/32      20600          20600          xe7         10.11.44.2
                                   20600          xe9         10.11.22.2
10.11.22.2/32       26881          3              xe9         10.11.22.2

DUT1#
```

Validate ECMP FTN, For the ECMP prefix 10.84.166.6/32, a single FTN entry is created with all the ECMP nexthops.

```
DUT1#show mpls forwarding-table 10.84.166.6/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Enabled
Code    FEC               FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-Intf
ELC       Nexthop         UpTime
   i>  10.84.166.6/32     5       40       -         (e)   -          -         -
-               00:23:09
                                     38       0           Yes   20600        xe7
No        10.11.44.2      -
                                     17       0           Yes   20600        xe9
No        10.11.22.2      -
DUT1#
```

Validate ECMP FTN, For the ECMP prefix 10.84.166.6/32, a single FTN entry is created with all the ECMP nexthops. For each nexthop, a cross-connect is created

```
DUT1#sh mpls ftn-table 10.84.166.6/32
 Primary FTN entry with FEC: 10.84.166.6/32, id: 5, row status: Active, Tunnel-Policy:
N/A, State: Installed
  CreateTime: 00:23:22, UpTime: 00:23:22, LastUpdate: N/A
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming
DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
     Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 38 refcount: 1
      Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 38, owner: ISIS-SR, Stale: NO, refcount: 6, out intf: xe7,
out label: 20600
    Nexthop addr: 10.11.44.2        cross connect ix: 8, op code: Push

     Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 17 refcount: 1
      Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
       Out-segment with ix: 17, owner: ISIS-SR, Stale: NO, refcount: 5, out intf: xe9,
out label: 20600
    Nexthop addr: 10.11.22.2        cross connect ix: 8, op code: Push

  Dependent service info (count 1):
  [CONFIRM_VRF] ftn_ix 1 owner BGP prefix 200.7.1.0/24 nhlfe_ix 51 vrf 2

DUT1#
```

Validation of SR-ECMP on DUT2[Transit Router]

This command displays the ILM-ID, FTN-ID, In-Label & Out-Label for all the IS-IS routes which have ILM/FTN entry installed. For the ECMP prefix 10.84.166.6//32, only one ILM entry will be installed, but Out-Label will be separate for each next-hop.

```
DUT1#sh ip isis route prefix 10.84.166.6/32 detail

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 1:  VRF : default
     Destination         Metric      ILM-ID     FTN-ID     In-Label   Next-Hop
Interface      Tag       Out-Label
ia   10.84.166.6/32      1040        9          5          20600      10.11.44.2
xe7            0         20600
                                                          10.11.22.2        xe9
0         20600
     Src: 0000.0001.0022 Ifindex 10016
     Src: 0000.0001.0044 Ifindex 10014
DUT1#
```

For the ECMP prefix 10.84.166.6/32, a single ILM entry is created with all the ECMP next-hops.

```
DUT2#sh mpls ilm-table 10.84.166.6/32
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
```

```
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

ILM-ECMP LDP: Disabled, SR: Enabled
Code    FEC/VRF/L2CKT    ILM-ID     In-Label    Out-Label   In-Intf    Out-Intf/VRF
Nexthop                  pri  UpTime
   i>   10.84.166.6/32     8          20600       20600       N/A        ce0
10.22.55.2               Yes  00:25:39
                                    20600       20600       N/A        xe8
10.22.33.2               Yes  -
DUT2#
```

CLI Commands

The Segment Routing ECMP feature introduces the following configuration commands.

mpls ilm-ecmp sr

Use this command to enable programming of SR ILM entry as ECMP in hardware. This command applies only to data-plane and IGP ECMP calculation does not depend on this CLI. Only if this command is enabled, SR ILM entry will be installed as ECMP entry in hardware with all the ECMP next-hops.

Use `no` parameter of this command to disable programming of SR ILM entry as ECMP in hardware. When `no` parameter of this command is executed, the installed SR ECMP ILM entry will be changed to SR non-ECMP ILM entry.

Command Syntax

```
mpls ilm-ecmp sr
```

```
no mpls ilm-ecmp sr
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced the `mpls ilm-ecmp sr` parameter in the OcNOS version 6.5.3.

Example

The following sequence of commands is used to enable programming of SR ILM entry as ECMP in hardware.

```
#configure terminal
```

```
(config)#mpls ilm-ecmp sr
```

```
(config)#
```

mpls ftn-ecmp sr

Use this command to enable programming of SR FTN entry as ECMP in hardware. This command applies only to data-plane and IGP ECMP calculation doesn't depend on this CLI. Only if this command is enabled, SR FTN entry will be installed as ECMP entry in hardware with all the ECMP nexthops.

Use `no` parameter of this command to disable programming of SR FTN entry as ECMP in hardware. When `no` parameter of this command is executed, the installed SR ECMP FTN entry will be changed to SR non-ECMP ILM entry.

Command Syntax

```
mpls ftn-ecmp sr
no mpls ftn-ecmp sr
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced the `mpls ftn-ecmp sr` parameter in the OcNOS version 6.5.3.

Example

The following sequence of commands is used to enable programming of SR FTN entry as ECMP in hardware.

```
#configure terminal
(config)#mpls ftn-ecmp sr
(config)#
```

show ip isis route prefix A.B.C.D/M

Use this command to display the ISIS routing table of the specified IPv4 prefix.

Command Syntax

```
show ip isis (WORD|) route ((prefix A.B.C.D/M)|)
```

Parameters

WORD	Information for a single ISIS area.
Prefix	Prefix.
A.B.C.D/M	IPv4 prefix.

Command Mode

Privileged exec mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

The following example displays the ISIS routing table of the specified IPv4 prefix.

```
#show ip isis route prefix 10.10.10.10/32

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 100:  VRF : default
     Destination         Metric      Next-Hop            Interface      Tag
L1   10.10.10.10/32      110         2.15.1.15           xe5            0
                                     2.10.1.10           xe11           0
#
```

show ip isis route detail

Use this command to display the MPLS information (ILM-ID, FTN-ID, In-label & Out-label) of the specified IS-IS IPv4 route or all the IS-IS IPv4 routes.

Command Syntax

```
show ip isis (WORD|) route ((prefix A.B.C.D/M)|) detail
```

Parameters

WORD	Information for a single ISIS area.
Prefix	Prefix.
A.B.C.D/M	IPv4 prefix.

Command Mode

Privileged exec mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

The following example displays the MPLS information of the specified prefix.

```
#show ip isis route prefix 10.10.10.10/32 detail

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 100:  VRF : default
     Destination         Metric      ILM-ID     FTN-ID     In-Label   Next-Hop
          Interface      Tag       Out-Label
L1   10.10.10.10/32      110         3          1          16010      2.15.1.15
          xe5            0         16010
                                                                      2.10.1.10
          xe11           0         16010
     Src: 0000.0000.0010 Ifindex 10011
     Src: 0000.0000.0010 Ifindex 10005
#

When the command is executed without prefix parameter, MPLS information of all the IS-IS
prefixes are displayed
#show ip isis route detail

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag 100:  VRF : default
     Destination         Metric      ILM-ID     FTN-ID     In-Label   Next-Hop
          Interface      Tag       Out-Label
C    2.2.2.2/32          10          1          --         16002      --
          lo             0         --
     Src: Connected IS-IS Interface
C    2.10.1.0/24         100         --         --         --         --
          xe11           0         --
     Src: Connected IS-IS Interface
C    2.15.1.0/24         50          --         --         --         --
```

```
          xe5             0         --
     Src: Connected IS-IS Interface
L1   10.10.10.10/32      110          3          1          16010      2.15.1.15
          xe5             0         16010
                                                                       2.10.1.10
          xe11            0         16010
     Src: 0000.0000.0010 Ifindex 10011
     Src: 0000.0000.0010 Ifindex 10005
L1   10.15.1.0/24        100          --         --         --         2.15.1.15
          xe5             0         --
     Src: 0000.0000.0015 Ifindex 10005
L1   15.15.15.15/32      60           5          2          16015      2.15.1.15
          xe5             0         16015
     Src: 0000.0000.0015 Ifindex 10005
#
```

show ip isis route tilfa prefix A.B.C.D/M

Use this command to display the MPLS information (SR outgoing label, PQ node, Backup outgoing label, Bypass trunk ID, Backup out-interface & Protection-type) of all ECMP next-hops of the specified IPv4 prefix. This is an enhancement to the existing command `show ip isis route tilfa` to insert `prefix` as an optional parameter.

Command Syntax

```
show ip isis route (WORD|) tilfa ((prefix A.B.C.D/M)|)
```

Parameters

`WORD`	Information for a single ISIS area.
`Prefix`	Prefix.
`A.B.C.D/M`	IPv4 prefix.

Command Mode

Privileged exec mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

```
#show ip isis route tilfa prefix 10.10.10.10/32

Tag   : 100  VRF : default
Codes : L1 - IS-IS level-1, L2 - IS-IS level-2,
```

```
          C - Connected Routes, ia - IS-IS inter area

10.10.10.10/32
   Route type: L1, FTN-ix :1  ILM-ix :3
   SR Incoming Label      : 16010
   Primary Path Nexthop   : 2.10.1.10, xe11
     SR outgoing Label    : 16010
     PQ node              : 15.15.15.15
     Backup outgoing Label: 16010
     Bypass_trunk id      : 2202
     Backup out interface : xe5
     Protection Type      : Link Protecting
   Primary Path Nexthop   : 2.15.1.15, xe5
     SR outgoing Label    : 16010
     PQ node              : 10.10.10.10
     Backup outgoing Label: 3
     Bypass_trunk id      : 2201
     Backup out interface : xe11
     Protection Type      : Node Protecting

   Trunk : 2201 :10.10.10.10_nh_10011_ALG0   FTN-ix : 3 ref_cnt:3
   Number Of outgoing label : 1
    16010
   Nexthop address : 2.10.1.10

#
```

show isis tilfa pq (WORD|)

Use this command to display the PQ nodes of all the ECMP next-hops of the specified vertex. This is an enhancement to the `isis tilfa pq` command to insert `system-id/hostname` as an optional parameter.

Command Syntax

```
show isis (WORD|) tilfa pq (WORD|)
```

Parameters

`WORD`	Information for a single ISIS area.
`WORD`	System-ID xxxx.xxxx.xxxx or hostname.

Command Mode

Privileged exec mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

When the command is executed by specifying `system-id` parameter:

```
#show isis tilfa pq 7010.00-00

Tag 100: Level-1 Link State Database:

Node: 7010.00-00
 Interface xe5
  P node: 0000.0000.0010 primary dist:100
  P node: 0000.0000.0015 primary dist:150
  Q node: 0000.0000.0010
  Q node: 0000.0000.0015
  Node Protecting P Nodes
  P node: 0000.0000.0010 primary dist:100

  PQ Node: 7010.00-00 backup dist:100
  PQ Node (Node Protection): 7010.00-00 backup dist:100
 Interface xe11
  P node: 0000.0000.0010 primary dist:100
  P node: 0000.0000.0015 primary dist:50
  Q node: 0000.0000.0010
  Q node: 0000.0000.0015
  Node Protecting P Nodes
  P node: 0000.0000.0015 primary dist:50

  PQ Node: 7015.00-00 backup dist:50
  No PQ Node found on backup path (Node Protection)
#
```

When the command is executed by specifying `hostname` parameter:

```
#show isis tilfa pq 7010

Tag 100: Level-1 Link State Database:

Node: 7010.00-00
 Interface xe5
```

```
  P node: 0000.0000.0010 primary dist:100
  P node: 0000.0000.0015 primary dist:150
  Q node: 0000.0000.0010
  Q node: 0000.0000.0015
  Node Protecting P Nodes
  P node: 0000.0000.0010 primary dist:100

  PQ Node: 7010.00-00 backup dist:100
  PQ Node (Node Protection): 7010.00-00 backup dist:100
 Interface xe11
  P node: 0000.0000.0010 primary dist:100
  P node: 0000.0000.0015 primary dist:50
  Q node: 0000.0000.0010
  Q node: 0000.0000.0015
  Node Protecting P Nodes
  P node: 0000.0000.0015 primary dist:50

  PQ Node: 7015.00-00 backup dist:50
  No PQ Node found on backup path (Node Protection)
#
```

show ip ospf route detail

Use this command to display the MPLS information (ILM-ID, FTN-ID, In-label & Out-label) of the specified OSPF IPv4 route or all the OSPF IPv4 routes.

Command Syntax

```
show ip ospf (<0-65535>|) route (((A.B.C.D | A.B.C.D/M |) detail) | )
```

Parameters

`<0-65535>`	Router process identifier.
`A.B.C.D`	Single route.
`A.B.C.D/M`	Single exact match route.

Command Mode

Privileged exec mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

The following example displays the MPLS information of all the OSPFv2 routes if a prefix parameter is not specified.

```
#show ip ospf route detail

OSPF process 100:
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

   Destination        Metric      Nexthop            Interface   Area
  ILM-ID     FTN-ID     In-Label   Out-Label
C  2.2.2.2/32         1           connected          lo          0.0.0.0
C  2.10.1.0/24        100         connected          xe11        0.0.0.0
C  2.15.1.0/24        50          connected          xe5         0.0.0.0
O  10.10.10.10/32    101       2.10.1.10        xe11       0.0.0.0          8
  5          19010      19010
                            2.15.1.15        xe5        0.0.0.0          8
  5          19010      19010
O  10.15.1.0/24       100         2.15.1.15          xe5         0.0.0.0
O  15.15.15.15/32    51        2.15.1.15        xe5        0.0.0.0          10
  6          19015      19015
#
```

The following example displays the MPLS information of the specified OSPFv2 route if a prefix parameter is specified:

```
#show ip ospf route 10.10.10.10/32 detail

OSPF process 100:
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       OSPF LFA attributes:
       P - Primary, SP - Secondary-Path, LP - Link Protecting,
       NP - Node Protecting, BID - Broadcast Link Protecting
       DP - Downstream Protecting

   Destination        Metric      Nexthop            Interface   Area
  ILM-ID     FTN-ID     In-Label   Out-Label
```

```
O  10.10.10.10/32    101       2.10.1.10        xe11       0.0.0.0          8
  5          19010      19010
                              2.15.1.15        xe5        0.0.0.0          8
  5          19010      19010
```

show hsl hw unit 0 encap-db LSP_ENCAP_ID

Use this command to display information on the `lsp_encap` entry installed in the hardware. This is an enhancement to the `show hsl hw unit 0 encap-db` command to insert `LSP_ENCAP_ID` as an optional parameter.

Command Syntax

```
show hsl hw unit 0 encap-db (LSP_ENCAP_ID|)
```

Parameters

`LSP_ENCAP_ID`	LSP Encap ID

Command Mode

Privileged exec mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Example

```
#show hsl hw unit 0 encap-db 0x40002044
label_array[0]:
Entropy enabled             : NO
flags                       : 8229(0x2025)
flags2                      : 0(0x0)
label                       : 16010
qos_map_id                  : 537133060
exp                         : 0
ttl                         : 64
pkt_pri                     : 0
pkt_cfi                     : 0
tunnel_id                   : 1073750084(0x40002044)
l3_intf_id                  : 4106(0x100a)
MPLS labelaction            : BCM_MPLS_EGRESS_ACTION_PUSH
egress_failover_id          : 0(0x0)
egress_failover_if_id       : 0(0x0)
outlif_counting_profile     : 0(0x0)
spl_label_push_type         : bcmMplsSpecialLabelPushNone
encap_access                :
estimated_encap_size        : 0(0x0)
```

```
#
```

Below are the revised commands. For more details, refer to the Segment Routing Commands chapter.

- show ip isis route tilfa
- show isis tilfa pq
- show hsl mpls tunnel (tunnel-id VALUE|)
- show ip ospf tilfa-backup-path
- show ip ospf tilfa-repair-list

Troubleshooting

1. If SR ILM entry is not installed as ECMP in hardware:
   - Check if `mpls ilm-ecmp sr` command is enabled.
   - Check if ECMP next-hops exist for that FEC in ISIS route table `show ip isis route detail` or OSPF route table `show ip ospf route detail`
2. If SR FTN entry is not installed as ECMP in hardware:
   - Check if `mpls ftn-ecmp sr` command is enabled.
   - Check if ECMP next-hops exist for the FEC in ISIS route table `show ip isis route detail` or OSPF route table `show ip ospf route detail`.
3. If traffic is not load-balanced among the ECMP next-hops:
   - Increase the number of flows, as load balancing depends on the internal hash computed by BCM. Note that lesser flows may lead to the same outgoing interface for different flows.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	**Description**
Label switched path (LSP)	A sequence of routers that co-operatively perform MPLS operations for a packet stream.
Topology-Independent Loop-Free Alternate (TI-LFA)	The ability to provide a loop free backup path irrespective of the topologies used in the network.
FEC-to-NHLFE (FTN)	A mapping from the forwarding equivalence class (FEC) of incoming packets to the corresponding Next Hop Label Forwarding Entry (NHLFE) in MPLS.
Incoming Label Map (ILM)	A mapping from incoming labels to corresponding Next Hop Label Forwarding Entry (NHLFE) in MPLS.
Interior Gateway Protocol (IGP)	An intra-domain protocol used to exchange network reachability and routing information among devices.
Forward Error Correction (FEC)	A system of error control that allows the receiver to correct some errors without having to request a re-transmission of data.

CHAPTER 2 ISIS Multi Topology

Overview

Intermediate System to Intermediate System (ISIS) is a link-state routing protocol commonly used in large-scale service provider networks and enterprise networks. By default, ISIS is in a single topology with no separate Shortest Path First (SPF) process to differentiate between IPv4 and IPv6 topologies. If the topology in IPv6 is different from IPv4, the routing calculation encounters a problem as the routes are evaluated and chosen based on the common topology.

Multi Topology (MT) is a mechanism to run a set of independent IP topologies within a single ISIS domain. This means, both IPv4 and IPv6 have different topologies in the network and two SPF processes are run to find the route to each IPv4 and IPv6 destination independently.

Feature Characteristics

The main characteristics of ISIS Multi Topology are as follows:

- Enables ISIS to maintain separate topologies for IPv4 and IPv6 within the same ISIS area or domain.
- Allows routers in the ISIS area (for Level 1 routing) or domain (for Level 2 routing) to support both IPv4 and IPv6 address families.
- Performs multiple SPF calculations for each configured topology.
- Defines new Type-Length-Value (TLV) encodings called Multi Topology TLV (MT TLV). It is used to advertise the multiple topologies supported by the routers and contains information about the topology, including the ID (MTID), flags, and MT metric.
  - MT TLV (229): Capability TLV advertised in Hello packets.
  - MT intermediate system TLV (222): Extended TLV that describes the adjacency between nodes once the adjacency is formed.
  - MT IPV6 reachability TLV (237): Reachability TLV that gives information on IPv6 routing.

Benefits

The key benefits of ISIS Multi Topology are as follows:

- Enables the ability to make changes to the IPv6 topology without affecting the IPv4 topology, and vice-versa.
- Leverages common adjacency and database tables.
- Provides an independent SPF process for IPv4 and IPv6.

Prerequisites

- To enable ISIS Multi Topology on OcNOS devices, wide metric configuration is mandatory.
- Follow the below configuration steps to prepare the interface for implementation of Multi Topology by enabling single topology on the routers:

  Note: In each of the commands, modify the relevant router as R1, R2, R3, R4 or R5, depending on the router being configured.

1. Enter configure mode followed by interface mode on loopback interface.
```
#configure terminal
R1(config)#int lo
```
2. Configure the IP address for the interface.
```
R1(config -if)# ip add 1.1.1.1/32 secondary
R1(config -if)# ipv6 address 1111::11/128
```
3. Include the interface in the router's ISIS 1 instance and exit the interface mode.
```
R1(config -if)# ip router isis 1
R1(config -if)# ipv6 router isis 1
R1(config -if)# exit
```
4. Enter the interface configuration mode and configure the IP address for the interface.
```
R1(config)#int xe22
R1(config -if)# ip address 10.1.1.1/24
R1(config -if)# ipv6 address 1001::1/64
```
5. Include the interface in the router's ISIS 1 instance and exit the interface mode.
```
R1(config -if)# ip router isis 1
R1(config -if)# ipv6 router isis 1
R1(config -if)# exit
```

For Routers R1 and R5, continue the configuration steps as follows:

6. Set the routing process ID as 1 and configure the IS type as level-1.
```
R1(config)# router isis 1
R1(config-router)# is-type level-1
```
7. Configure wide metric-style.
```
R1(config-router)# metric-style wide
```
8. Enable dynamic host name under ISIS process.
```
R1(config-router)# dynamic-hostname
```
9. Enable BFD in all the interfaces.
```
R1(config-router)# bfd all-interfaces
```
10. Configure Network Entity Title (NET).
```
R1(config-router)# net 49.0000.0000.0001.00
```
11. Commit the candidate configuration to the running configuration.
```
R1(config-router)# commit
```

For Routers R2 and R4, use the following configuration steps after you exit the interface mode (step 5 shown above):

1. Enter the interface configuration mode and configure the IP address for the interface.
```
R2(config)#int xe24
R2(config -if)# ip address 20.1.1.1/24
R2(config -if)# ipv6 address 2001::1/64
```
2. Include the interface in the router's ISIS 1 instance and exit the interface mode.
```
R2(config -if)# ipv6 router isis 1
R2(config -if)# exit
```
3. Enter the interface configuration mode and configure the IP address for the interface.
```
R2(config)#int xe23
```

```
R2(config -if)# ip address 40.1.1.1/24
R2(config -if)# ipv6 address 4001::1/64
```

4. Include the interface in the router's ISIS 1 instance and exit the interface mode.

```
R2(config -if)# ip router isis 1
R2(config -if)# ipv6 router isis 1
R2(config -if)# exit
```

5. Set the routing process ID as 1 and configure IS type as level 1.

```
R2(config)# router isis 1
R2(config-router)# is-type level-1
```

6. Configure wide metric style.

```
R2(config-router)# metric-style wide
```

7. Enable dynamic host name under ISIS process.

```
R2(config-router)# dynamic-hostname
```

8. Enable BFD in all the interfaces.

```
R2(config-router)# bfd all-interfaces
```

9. Configure Network Entity Title (NET).

```
R2(config-router)# net 49.0000.0000.0002.00
```

10. Commit the candidate configuration to the running configuration.

```
R2(config-router)# commit
```

For Router R3, follow these configuration steps after you exit the interface mode (step 5 shown above):

1. Enter Interface configuration mode and configure the IP address of the interface.

```
R3(config)#int xe31/1
R3(config -if)# ip address 50.1.1.1/24
R3(config -if)# ipv6 address 5001::1/64
```

2. Include the interface in the router's ISIS 1 instance and exit the interface mode.

```
R3(config -if)# ip router isis 1
R3(config -if)# ipv6 router isis 1
R3(config -if)# exit
```

3. Set the routing process ID as 1 and configure IS type as level-1.

```
R3(config)# router isis 1
R3(config-router)# is-type level-1
```

4. Configure wide metric-style.

```
R3(config-router)# metric-style wide
```

5. Enable dynamic host name under ISIS process.

```
R3(config-router)# dynamic-hostname
```

6. Enable BFD on all the interfaces.

```
R3(config-router)# bfd all-interfaces
```

7. Configure Network Entity Title (NET).

```
R3(config-router)# net 49.0000.0000.0003.00
```

8. Commit the candidate configuration to the running configuration.

```
R3(config-router)# commit
```

Configuration

To set up Multi Topology in ISIS, the configuration is as shown below:

Topology

This topology diagram consists of five routers (R1, R2,R3,R4 and R5).

It has both ISIS IPv4 and IPv6 routing enabled, except the link between R2 and R4 which has only IPv6 enabled.

In Single Topology, router R1 receives the information and calculates a SPF tree. To reach 5.5.5.5 (R5 IPv4), it takes the path R1-> R2 -> R4 ->R5. However, it fails since R2 to R4 is solely an IPv6 path. Since the same SPF tree is used for both IPv4 and IPv6 in R1, it considers the link between R2 -> R4 as the shortest path instead of R2 -> R3 -> R4.

On enabling Multi Topology on all the routers, SPF trees are calculated separately for IPv4 and IPv6 routing. This means, to reach from R1 to R5, IPv4 takes the R1 -> R2 -> R3 -> R4 -> R5 path and IPv6 takes the R1 -> R2 -> R4 -> R5 path.

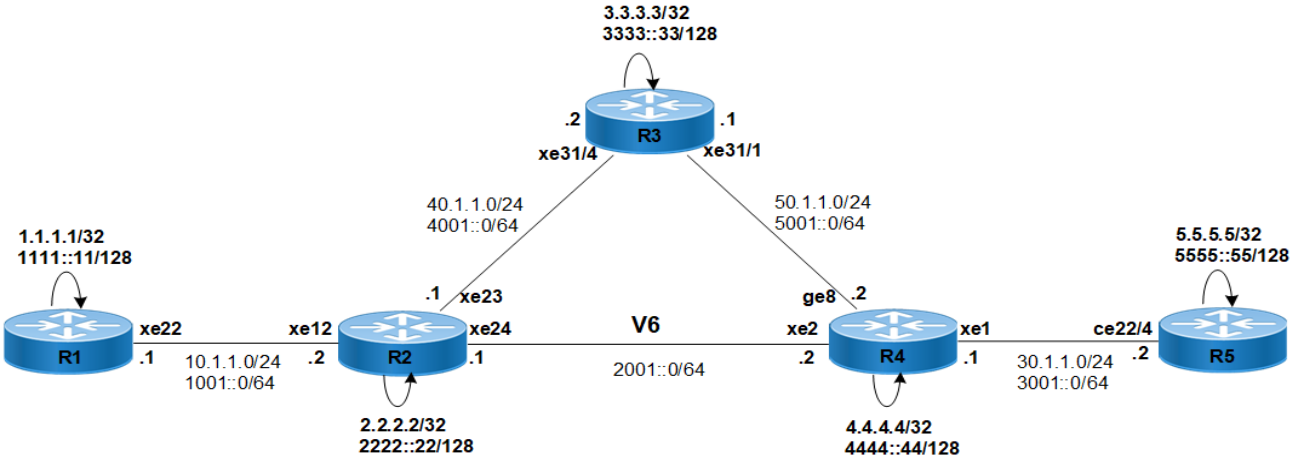

Figure 2-15: ISIS Multi Topology

To configure multi topology on the routers R1, R2, R3, R4 and R5, follow the steps mentioned below:

Note: Ensure that the Prerequisites are met for all the routers.

Note: Modify the commands for the relevant routers being configured (R1, R2, R3, R4 or R5).

1. Set the routing process ID as 1.

```
R1(config)# router isis 1
```

2. Configure metric-style wide.

```
R1(config-router)# metric-style wide
```

3. Configure address family IPv6.

```
R1(config-router)#address-family ipv6
```

4. Enable multi topology with level 1.

```
R1(config-router-af)#multi-topology level-1
```

5. Commit the candidate configuration to the running configuration.

```
R1(config-router-af)#commit
```

Validation for Multi Topology

```
R1#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol
R2             xe22        00e0.4b77.39fe      Up     19        L1   M-ISIS

R1#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R2             xe22        Up     L1   64        0000.0000.0001.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:09:39
  Area Address(es): 49
  IP Address(es): 10.1.1.2
  IPv6 Address(es): fe80::2e0:4bff:fe77:39fe
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise


R1#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
R2                   10         R2       xe22        00e0.4b77.39fe
R3                   20         R2       xe22        00e0.4b77.39fe
R4                   30         R2       xe22        00e0.4b77.39fe
R5                   40         R2       xe22        00e0.4b77.39fe


R1#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   --
```

```
R2                  10         R2       xe22        00e0.4b77.39fe
R3                  20         R2       xe22        00e0.4b77.39fe
R4                  20         R2       xe22        00e0.4b77.39fe
R5                  30         R2       xe22        00e0.4b77.39fe


R1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, installed 01:55:53, last update
01:55:53 ago
i L1         2.2.2.2/32 [115/20] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1         3.3.3.3/32 [115/30] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1         4.4.4.4/32 [115/40] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
i L1         5.5.5.5/32 [115/50] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
C            10.1.1.0/24 is directly connected, xe22, installed 01:55:53, last update
01:55:53 ago
i L1         30.1.1.0/24 [115/40] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
i L1         40.1.1.0/24 [115/20] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1         50.1.1.0/24 [115/30] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
C            127.0.0.0/8 is directly connected, lo, installed 01:57:14, last update
01:57:14 ago

Gateway of last resort is not set

R1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 01:57:15, last update 01:57:15 ago
C       1001::/64 via ::, xe22, installed 01:32:33, last update 01:32:33 ago
C       1111::11/128 via ::, lo, installed 01:33:09, last update 01:33:09 ago
```

```
i L1    2001::/64 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    2222::22/128 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1    3001::/64 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    3333::33/128 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1    4001::/64 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    4444::44/128 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1    5001::/64 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1    5555::55/128 [115/40] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
C       fe80::/64 via ::, xe25, installed 01:56:18, last update 01:56:18 ago

R1#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:09:57.608 ago


R1#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000015   0x9E64        602               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00            * 0x0000000C   0x724E        602               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        601               0/0/0
  Area Address: 49
```

```
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        601               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        601               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        601               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
```

```
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        601               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        601               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        601               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        601               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R1#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00            * 0x00000015   0x9E64        596               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
```

```
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00            * 0x0000000C   0x724E        596               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        595               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
```

```
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        595               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        595               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        595               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        595               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R2:

R2#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R1             xe12        e8c5.7a69.446f  Up     6         L1   M-ISIS
```

```
R3              xe23        903c.b3c5.ae9b      Up     6         L1   M-ISIS
R4              xe24        9819.2ccf.ede3      Up     9         L1   M-ISIS

R2#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R1             xe12        Up     L1   64        0000.0000.0001.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:10:56
  Area Address(es): 49
  IP Address(es): 10.1.1.1
  IPv6 Address(es): fe80::eac5:7aff:fe69:446f
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R3             xe23        Up     L1   64        0000.0000.0003.03
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:10:56
  Area Address(es): 49
  IP Address(es): 40.1.1.2
  IPv6 Address(es): fe80::923c:b3ff:fec5:ae9b
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R4             xe24        Up     L1   64        0000.0000.0004.04
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:10:56
  Area Address(es): 49
  IPv6 Address(es): fe80::9a19:2cff:fecf:ede3
  Topology: IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise


R2#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id           Metric     Next-Hop            Interface   SNPA
R1                  10         R1       xe12       e8c5.7a69.446f
R2                  --
```

```
R3                  10         R3       xe23        903c.b3c5.ae9b
R4                  20         R3       xe23        903c.b3c5.ae9b
R5                  30         R3       xe23        903c.b3c5.ae9b


R2#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id           Metric     Next-Hop             Interface   SNPA
R1                  10         R1       xe12        e8c5.7a69.446f
R2                  --
R3                  10         R3       xe23        903c.b3c5.ae9b
R4                  10         R4       xe24        9819.2ccf.ede3
R5                  20         R4       xe24        9819.2ccf.ede3

R2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/20] via 10.1.1.1, xe12, installed 01:11:03, last update
01:11:03 ago
C            2.2.2.2/32 is directly connected, lo, installed 01:59:20, last update
01:59:20 ago
i L1         3.3.3.3/32 [115/20] via 40.1.1.2, xe23, installed 01:11:03, last update
01:11:03 ago
i L1         4.4.4.4/32 [115/30] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
i L1         5.5.5.5/32 [115/40] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
C            10.1.1.0/24 is directly connected, xe12, installed 01:57:30, last update
01:57:30 ago
C            20.1.1.0/24 is directly connected, xe24, installed 01:59:19, last update
01:59:19 ago
i L1         30.1.1.0/24 [115/30] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
C            40.1.1.0/24 is directly connected, xe23, installed 01:59:19, last update
01:59:19 ago
i L1         50.1.1.0/24 [115/20] via 40.1.1.2, xe23, installed 01:11:03, last update
01:11:03 ago
C            127.0.0.0/8 is directly connected, lo, installed 02:20:04, last update
02:20:04 ago

Gateway of last resort is not set

R2#show ipv6 route
```

```
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:20:05, last update 02:20:05 ago
C       1001::/64 via ::, xe12, installed 01:32:42, last update 01:32:42 ago
i L1    1111::11/128 [115/20] via fe80::eac5:7aff:fe69:446f, xe12, installed 00:11:04,
last update 00:11:04 ago
C       2001::/64 via ::, xe24, installed 01:59:20, last update 01:59:20 ago
C       2222::22/128 via ::, lo, installed 01:33:21, last update 01:33:21 ago
i L1   3001::/64 [115/20] via fe80::9a19:2cff:fecf:ede3, xe24, installed 00:11:04, last
update 00:11:04 ago
i L1    3333::33/128 [115/20] via fe80::923c:b3ff:fec5:ae9b, xe23, installed 01:11:04,
last update 01:11:04 ago
C       4001::/64 via ::, xe23, installed 01:24:52, last update 01:24:52 ago
i L1    4444::44/128 [115/20] via fe80::9a19:2cff:fecf:ede3, xe24, installed 00:11:04,
last update 00:11:04 ago
i L1   5001::/64 [115/20] via fe80::923c:b3ff:fec5:ae9b, xe23, installed 01:11:04, last
update 00:11:04 ago
                  [115/20] via fe80::9a19:2cff:fecf:ede3, xe24
i L1    5555::55/128 [115/30] via fe80::9a19:2cff:fecf:ede3, xe24, installed 00:11:04,
last update 00:11:04 ago
C       fe80::/64 via ::, xe12, installed 01:57:31, last update 01:57:31 ago


R2#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:11:11.544 ago



R2#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
```

```
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        527               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00            * 0x00000014   0x2A52        528               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
```

```
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        527               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        527               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        527               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        527               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        527               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00
```

```
R2#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        520               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00            * 0x00000014   0x2A52        521               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
```

```
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        520               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        520               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        520               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        520               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        520               0/0/0
  Metric:   0          IS-Extended R5.00
```

```
  Metric:   0          IS-Extended R4.00


R3:



R3#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R4             xe31/1      9819.2ccf.ede9      Up     9         L1   M-ISIS
R2             xe31/4      00e0.4b77.3a09      Up     27        L1   M-ISIS


R3#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R4             xe31/1      Up     L1   64        0000.0000.0004.01
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:11:42
  Area Address(es): 49
  IP Address(es): 50.1.1.2
  IPv6 Address(es): fe80::9a19:2cff:fecf:ede9
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
R2             xe31/4      Up     L1   64        0000.0000.0003.03
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:11:42
  Area Address(es): 49
  IP Address(es): 40.1.1.1
  IPv6 Address(es): fe80::2e0:4bff:fe77:3a09
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise


R3#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
```

```
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       xe31/4      00e0.4b77.3a09
R2                   10         R2       xe31/4      00e0.4b77.3a09
R3                   --
R4                   10         R4       xe31/1      9819.2ccf.ede9
R5                   20         R4       xe31/1      9819.2ccf.ede9


R3#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       xe31/4      00e0.4b77.3a09
R2                   10         R2       xe31/4      00e0.4b77.3a09
R3                   --
R4                   10         R4       xe31/1      9819.2ccf.ede9
R5                   20         R4       xe31/1      9819.2ccf.ede9

R3#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/30] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
i L1         2.2.2.2/32 [115/20] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
C            3.3.3.3/32 is directly connected, lo, installed 02:00:27, last update
02:00:27 ago
i L1         4.4.4.4/32 [115/20] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
i L1         5.5.5.5/32 [115/30] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
i L1        10.1.1.0/24 [115/20] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
i L1        30.1.1.0/24 [115/20] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
C           40.1.1.0/24 is directly connected, xe31/4, installed 02:00:09, last update
02:00:09 ago
C           50.1.1.0/24 is directly connected, xe31/1, installed 02:00:26, last update
02:00:26 ago
C            127.0.0.0/8 is directly connected, lo, installed 02:18:52, last update
02:18:52 ago

Gateway of last resort is not set
```

```
R3#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:18:53, last update 02:18:53 ago
i L1    1001::/64 [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1    1111::11/128 [115/30] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1    2001::/64 [115/20] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
                  [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4
i L1    2222::22/128 [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1    3001::/64 [115/20] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C       3333::33/128 via ::, lo, installed 01:31:50, last update 01:31:50 ago
C       4001::/64 via ::, xe31/4, installed 01:30:10, last update 01:30:10 ago
i L1    4444::44/128 [115/20] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C       5001::/64 via ::, xe31/1, installed 01:29:43, last update 01:29:43 ago
i L1    5555::55/128 [115/30] via fe80::9a19:2cff:fecf:ede9, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C       fe80::/64 via ::, xe31/4, installed 02:00:10, last update 02:00:10 ago


R3#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:12:00.519 ago



R3#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
```

```
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        478               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00            * 0x00000013   0x7FCC        479               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
```

```
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00            * 0x0000000C   0x6D4E        479               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        478               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        478               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        478               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        478               0/0/0
  Metric:   0          IS-Extended R5.00
```

```
  Metric:   0          IS-Extended R4.00

R3#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        471               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00            * 0x00000013   0x7FCC        472               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
```

```
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00           * 0x0000000C   0x6D4E        472               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00             0x00000015   0x8C0D        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00             0x00000007   0x9A25        471               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00             0x0000000C   0x6751        471               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00             0x00000010   0xFA0F        471               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00             0x00000007   0xA813        471               0/0/0
```

```
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00



R4:

R4#show clns neighbors

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
R5             xe1         e001.a6aa.0f23  Up     6         L1   M-ISIS
R2             xe2         00e0.4b77.3a0a  Up     22        L1   M-ISIS
R3             ge8         903c.b3c5.ae98  Up     22        L1   M-ISIS


R4#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R5             xe1         Up     L1   64        0000.0000.0005.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:12:38
  Area Address(es): 49
  IP Address(es): 30.1.1.2
  IPv6 Address(es): fe80::e201:a6ff:feaa:f23
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R2             xe2         Up     L1   64        0000.0000.0004.04
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:12:37
  Area Address(es): 49
  IPv6 Address(es): fe80::2e0:4bff:fe77:3a0a
  Topology: IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R3             ge8         Up     L1   64        0000.0000.0004.01
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:12:38
```

```
  Area Address(es): 49
  IP Address(es): 50.1.1.1
  IPv6 Address(es): fe80::923c:b3ff:fec5:ae98
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

R4#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   30         R3       ge8         903c.b3c5.ae98
R2                   20         R3       ge8         903c.b3c5.ae98
R3                   10         R3       ge8         903c.b3c5.ae98
R4                   --
R5                   10         R5       xe1         e001.a6aa.0f23


R4#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   20         R2       xe2         00e0.4b77.3a0a
R2                   10         R2       xe2         00e0.4b77.3a0a
R3                   10         R3       ge8         903c.b3c5.ae98
R4                   --
R5                   10         R5       xe1         e001.a6aa.0f23


R4#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/40] via 50.1.1.1, ge8, installed 00:12:48, last update
00:12:48 ago
i L1         2.2.2.2/32 [115/30] via 50.1.1.1, ge8, installed 00:12:48, last update
00:12:48 ago
i L1         3.3.3.3/32 [115/20] via 50.1.1.1, ge8, installed 01:01:13, last update
01:01:13 ago
C            4.4.4.4/32 is directly connected, lo, installed 02:01:55, last update
02:01:55 ago
```

```
i L1         5.5.5.5/32 [115/20] via 30.1.1.2, xe1, installed 01:12:47, last update
01:12:47 ago
i L1         10.1.1.0/24 [115/30] via 50.1.1.1, ge8, installed 00:12:48, last update
00:12:48 ago
C            20.1.1.0/24 is directly connected, xe2, installed 02:01:04, last update
02:01:04 ago
C            30.1.1.0/24 is directly connected, xe1, installed 02:01:55, last update
02:01:55 ago
i L1         40.1.1.0/24 [115/20] via 50.1.1.1, ge8, installed 01:01:13, last update
01:01:13 ago
C            50.1.1.0/24 is directly connected, ge8, installed 02:01:22, last update
02:01:22 ago
C            127.0.0.0/8 is directly connected, lo, installed 02:20:17, last update
02:20:17 ago

Gateway of last resort is not set


R4#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:20:18, last update 02:20:18 ago
i L1    1001::/64 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48, last
update 00:12:48 ago
i L1    1111::11/128 [115/30] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48,
last update 00:12:48 ago
C       2001::/64 via ::, xe2, installed 02:01:05, last update 02:01:05 ago
i L1    2222::22/128 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48,
last update 00:12:48 ago
C       3001::/64 via ::, xe1, installed 01:33:20, last update 01:33:20 ago
i L1    3333::33/128 [115/20] via fe80::923c:b3ff:fec5:ae98, ge8, installed 01:01:14,
last update 01:01:14 ago
i L1    4001::/64 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 01:04:04, last
update 00:12:48 ago
                  [115/20] via fe80::923c:b3ff:fec5:ae98, ge8
C       4444::44/128 via ::, lo, installed 01:33:04, last update 01:33:04 ago
C       5001::/64 via ::, ge8, installed 01:29:27, last update 01:29:27 ago
i L1    5555::55/128 [115/20] via fe80::e201:a6ff:feaa:f23, xe1, installed 00:12:48,
last update 00:12:48 ago
C       fe80::/64 via ::, xe2, installed 02:01:05, last update 02:01:05 ago


R4#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
```

```
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:12:55.361 ago



R4#show isis database verbose
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        423               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        423               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        423               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
```

```
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        423               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        423               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00            * 0x00000015   0x8C0D        424               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00            * 0x00000007   0x9A25        424               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00            * 0x0000000C   0x6751        424               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        423               0/0/0
```

```
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        423               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R4#show isis database detail
Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        417               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
```

```
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
  Metric:   10         IP-Extended 10.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        417               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00            * 0x00000015   0x8C0D        418               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 4.4.4.4/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00            * 0x00000007   0x9A25        418               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
```

```
R4.04-00            * 0x0000000C   0x6751        418               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00              0x00000010   0xFA0F        417               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
  Metric:   10         IP-Extended 30.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00              0x00000007   0xA813        417               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00


R5:

R5#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1:  VRF : default
System Id      Interface   SNPA                State  Holdtime  Type Protocol
R4             ce22/4      9819.2ccf.ede2      Up     28        L1   M-ISIS


R5#show clns is-neighbors detail

Tag 1:  VRF : default
System Id      Interface   State  Type Priority  Circuit Id
R4             ce22/4      Up     L1   64        0000.0000.0005.02
  L1  Adjacency ID: 1
  L2  Adjacency ID: 2
  Uptime: 01:13:32
  Area Address(es): 49
  IP Address(es): 30.1.1.1
  IPv6 Address(es): fe80::9a19:2cff:fecf:ede2
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
```

```
R5#show isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   40         R4       ce22/4      9819.2ccf.ede2
R2                   30         R4       ce22/4      9819.2ccf.ede2
R3                   20         R4       ce22/4      9819.2ccf.ede2
R4                   10         R4       ce22/4      9819.2ccf.ede2
R5                   --


R5#show ipv6 isis topology

Tag 1:  VRF : default
IS-IS paths to level-1 routers
System Id            Metric     Next-Hop             Interface   SNPA
R1                   30         R4       ce22/4      9819.2ccf.ede2
R2                   20         R4       ce22/4      9819.2ccf.ede2
R3                   20         R4       ce22/4      9819.2ccf.ede2
R4                   10         R4       ce22/4      9819.2ccf.ede2
R5                   --

R5#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1         1.1.1.1/32 [115/50] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
i L1         2.2.2.2/32 [115/40] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
i L1         3.3.3.3/32 [115/30] via 30.1.1.1, ce22/4, installed 01:02:05, last update
01:02:05 ago
i L1         4.4.4.4/32 [115/20] via 30.1.1.1, ce22/4, installed 01:13:40, last update
01:13:40 ago
C            5.5.5.5/32 is directly connected, lo, installed 02:03:15, last update
02:03:15 ago
i L1        10.1.1.0/24 [115/40] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
C           30.1.1.0/24 is directly connected, ce22/4, installed 02:03:15, last update
02:03:15 ago
i L1        40.1.1.0/24 [115/30] via 30.1.1.1, ce22/4, installed 01:04:55, last update
01:04:55 ago
i L1        50.1.1.0/24 [115/20] via 30.1.1.1, ce22/4, installed 01:02:05, last update
01:02:05 ago
```

```
C            127.0.0.0/8 is directly connected, lo, installed 02:20:59, last update
02:20:59 ago

Gateway of last resort is not set


R5#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 P - SRV6-POLICY,
        v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, installed 02:21:00, last update 02:21:00 ago
i L1    1001::/64 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    1111::11/128 [115/40] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    2001::/64 [115/20] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    2222::22/128 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
C       3001::/64 via ::, ce22/4, installed 01:05:32, last update 01:05:32 ago
i L1    3333::33/128 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    4001::/64 [115/30] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    4444::44/128 [115/20] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1    5001::/64 [115/20] via fe80::9a19:2cff:fecf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
C       5555::55/128 via ::, lo, installed 01:06:20, last update 01:06:20 ago
C       fe80::/64 via ::, ce22/4, installed 02:03:16, last update 02:03:16 ago


R5#show isis spf-logs level-1-2
Tag 1:  VRF : default
Level-1 spf logs:
 Next SPF is not scheduled yet
 SPF schedule delay min 0 secs 500 msecs
 SPF schedule delay max 50 secs 0 msecs
 SPF algorithm executed 12 times
 SPF algorithm last executed 00:13:45.938 ago



R5#show isis database verbose
Tag 1:  VRF : default
```

```
IS-IS Level-1 Link State Database:
LSPID                 LSP Seq Num  LSP Checksum  LSP Holdtime      ATT/P/OL
R1.00-00              0x00000015   0x9E64        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 1111::11/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
R1.02-00              0x0000000C   0x724E        373               0/0/0
  Metric:   0          IS-Extended R1.00
  Metric:   0          IS-Extended R2.00
R2.00-00              0x00000014   0x2A52        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:   10         IS-Extended R1.02
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R1.02
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 40.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 2222::22/128
  Metric:   10         IPv6 (MT-IPv6) 1001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
R3.00-00              0x00000013   0x7FCC        372               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric:   10         IS-Extended R4.01
```

```
  Metric:   10         IS-Extended R3.03
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IS (MT-IPv6) R3.03
  Metric:   10         IP-Extended 3.3.3.3/32
  Metric:   10         IP-Extended 50.1.1.0/24
  Metric:   10         IP-Extended 40.1.1.0/24
  Metric:   10         IPv6 (MT-IPv6) 3333::33/128
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
  Metric:   10         IPv6 (MT-IPv6) 4001::/64
R3.03-00              0x0000000C   0x6D4E        372               0/0/0
  Metric:   0          IS-Extended R3.00
  Metric:   0          IS-Extended R2.00
R4.00-00              0x00000015   0x8C0D        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS-Extended R4.01
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IS (MT-IPv6) R4.04
  Metric:   10         IS (MT-IPv6) R4.01
  Metric:   10         IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 4444::44/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
  Metric:   10         IPv6 (MT-IPv6) 2001::/64
  Metric:   10         IPv6 (MT-IPv6) 5001::/64
R4.01-00              0x00000007   0x9A25        373               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R3.00
R4.04-00              0x0000000C   0x6751        373               0/0/0
  Metric:   0          IS-Extended R4.00
  Metric:   0          IS-Extended R2.00
R5.00-00            * 0x00000010   0xFA0F        373               0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric:   10         IS-Extended R5.02
  Metric:   10         IS (MT-IPv6) R5.02
  Metric:   10         IP-Extended 5.5.5.5/32
```

```
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:   10         IPv6 (MT-IPv6) 5555::55/128
  Metric:   10         IPv6 (MT-IPv6) 3001::/64
R5.02-00            * 0x00000007   0xA813        373               0/0/0
  Metric:   0          IS-Extended R5.00
  Metric:   0          IS-Extended R4.00
```

Running Configuration

```
R1#sh running-config router isis
!
router isis 1
 is-type level-1
 metric-style wide
 dynamic-hostname
 bfd
 all-interfaces
 net 49.0000.0000.0001.00
 !
 address-family ipv6
 multi-topology
 level-1
 exit-address-family
!
R1#
```

CLI Commands

The ISIS Multi-topology feature introduces the `multi-topology` configuration command.

multi topology

Use this command to configure the ISIS topology type.

Use `no` parameter of this command to set the topology back to single.

Command Syntax

```
multi-topology (level-1|level-1-2|level-2)
no multi-topology
```

Parameters

`level-1`	Specify to enable multi-topology for level 1
`level-2`	Specify to enable multi-topology for level 2
`level-1-2`	Specify to enable multi-topology for both the levels

Default

ISIS topology type applies to levels 1 and 2.

Command Mode

Address-family IPv6 mode.

Applicability

Introduced the `multi-topology` parameter in OcNOS version 6.5.2.

Example

The following sequence of commands is used to configure ISIS `multi-topology` type as transition for levels 1 and 2.

```
(config)#router isis 1
(config-router)#address-family ipv6 unicast
(config-router-af)#multi-topology level-1-2
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
ISIS	Intermediate System to Intermediate System is a link-state routing protocol.
Multi Topology (MT)	In ISIS, Multi Topology (MT) is a mechanism to run a set of independent IP topologies within a single ISIS domain.
Type Length Value (TLV)	A data structure used to encode optional information in a data communications protocol: • Type: the kind of field that this part of the message represents • Length: the size of the value field, usually in bytes • Value: a variable-sized set of bytes that contains the data of the message
Shortest Path First (SPF)	Algorithm used by ISIS to make routing decisions based on the state of network links.
Loopback	A troubleshooting test in which a signal is transmitted from a source to a destination and then back to the source again so that the signal can be measured and evaluated.
Wide metric configuration	Allows ISIS to support larger networks by configuring high value metric in the interface.
Hello Packets	Information packets used to discover ISIS neighbors and maintain adjacencies.
Link State Packets (LSP)	Unidirectional, point-to-point, half-duplex connection used to exchange link state information.